
University of Alberta

Hypertext navigation tools as mechanisms for the investigation of hyperspace properties: Spatial and conceptual relations, metric space and mental representation.

by

Patricia M. Boechler ©

A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment
of the requirements for the degree of Doctor of Philosophy

Department of Psychology

Edmonton, Alberta

Fall, 2002

National Library
of Canada

Acquisitions and
Bibliographic Services

395 Wellington Street
Ottawa ON K1A 0N4
Canada

Bibliothèque nationale
du Canada

Acquisitions et
services bibliographiques

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence

Our file Notre référence

The author has granted a non-exclusive licence allowing the National Library of Canada to reproduce, loan, distribute or sell copies of this thesis in microform, paper or electronic formats.

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque nationale du Canada de reproduire, prêter, distribuer ou vendre des copies de cette thèse sous la forme de microfiche/film, de reproduction sur papier ou sur format électronique.

The author retains ownership of the copyright in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

L'auteur conserve la propriété du droit d'auteur qui protège cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

0-612-81166-2

Canada

University of Alberta

Library Release Form

Name of Author: Patricia Martha Boechler

Title of Thesis: Hypertext navigation tools as mechanisms for the investigation of hyperspace properties: Spatial and conceptual relations, metric space and mental representation.

Degree: Doctor of Philosophy

Year this Degree Granted: 2002

Permission is hereby granted to the university of Alberta Library to reproduce single copies of this thesis and to lend or sell such copies for private, scholarly or scientific research purposes only.

The author reserves all other rights in association with the copyright in this thesis, and except as herein before provided, neither the thesis nor any substantial portion thereof may be printed or otherwise reproduced in any material form whatever without the author's prior written permission.

P Boechler

17918-63A Ave.
Edmonton, Alberta, Canada
T5T 2J2

Date: August 14 / 2002

University of Alberta

Faculty of Graduate Studies and Research

The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies and Research for acceptance, a thesis entitled: "Hypertext navigation tools as mechanisms for the investigation of hyperspace properties: Spatial and conceptual relations, metric space and mental representation" submitted by Patricia M. Boechler in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

Dr. Michael R.W. Dawson
(Supervisor)

Dr. Edward Cornell

Dr. Michael Carbonaro

Dr. Wesley Cooper

Dr. Saul Greenberg
(External Examiner)

Date: AUGUST 08/02

ABSTRACT

This thesis is divided into two distinct but related inquiries. First, to investigate the impact of different information in hypertext navigation tools, four groups of users were each given a different navigation tool to complete an information search task through a 22-page hypertext document. Each tool contained a preponderance of spatial or hierarchical information, a combination of both or the absence of both. Measures were recorded of the overall task time, number of pages accessed and memory measures for the page titles in the document. As well, a distance-like ratings task was employed to tap into users' mental representations of the document. The results indicated that providing spatial or hierarchical or spatial/hierarchical information combined improves users' abilities to efficiently move through a document but does not enhance recall of page titles over that of users given an alphabetical list. The ratings task indicated that mental representations of the document were comparable across all four groups. However, relationships between the representations and the actual path patterns of users were only evident for users given spatial information.

The second line of inquiry in this thesis focuses on investigating behaviour/representation discrepancies through the use of artificial neural networks. An in-depth analysis of two neural networks trained to complete a spatial task revealed that different underlying representations can produce similar behaviours. In this case, a metric and a non-metric representation both produced behaviours that adhered to metric principles. A coarse coding scheme centred around virtual, optimal locations for each of

the processing units in the networks appears responsible for the network's ability to process both types of information and produce accurate spatial judgments.

Acknowledgements

This thesis would not have been possible without the guidance of my supervisor, Dr. Michael Dawson, a scholar whose breadth of knowledge, insightfulness and curiosity for science motivated me by his example.

I would also like to express my appreciation to a number of other people who contributed to the completion of this thesis:

To the members of the Biological Computation Project, I could not have asked for a better cohort with which to share this journey.

To Dorothy Steffler, for her caring advice and friendship throughout my graduate studies.

To MJ, whose unique combination of creativity, meticulousness and tenacity contributed much to this project.

To Theresa Wilson, for listening.

To my parents, Martha McAvoy and Robert Boechler, for reminding me often of the value of pursuing this goal and their certainty that I could achieve it.

I thank my siblings, Kelvin Boechler and Terri Blue for the endless hours of encouragement and constructive feedback offered at every turn.

Finally, I would like to thank my children, Camryn and Jace, for their unwavering patience and faith, and for their sacrifices throughout this endeavor. Beginning to end, you are my reason for the work I pursue.

Table of Contents

LIST OF TABLES
LIST OF FIGURES

CHAPTER 1 - INTRODUCTION

1.1 A Roadmap 2
1.2 An introduction to Hypertext 3
    1.2.1 Goals in Hypertext Use 5
    1.2.2 Advantages and Disadvantages of Hypertext 6
        Cognitive Overhead 7
        Disorientation 8
1.3 The Spatial Metaphor 9
    1.3.1 The Role of Metaphor 12
    1.3.2 Mental Models vs. Mental Representations 16
    1.3.3 Relating Space to Hypertext 17
        Metric Models of Space 17
        Spatial Primitives 22
1.4 Spatial Cognition 25
    1.4.1 Cognitive Maps 26
    1.4.2 Spatial Knowledge Acquisition 29
    1.4.3 A Preliminary Comparison of Hyperspace and Physical Space Features 32
        Perceptual Cues 32
        The Nature of Landmarks 33
        Spatial Mental Representations 34
        A Question of Scale 34
1.5 Hypertext Navigation 35
    1.5.1 Measures 36
    1.5.2 Types of Navigational Aids 37
    1.5.3 The Effects of Navigational Aids 38
        The Effects of Navigational Aids on Closed Tasks 39
        The Effects of Navigational Aids on Open Tasks 42
1.6 Summary of Review 44
1.7 The Current Research 44

CHAPTER 2 - TRACKING COGNITIVE PROCESSES IN HYPERTEXT NAVIGATION

2.1 The Data Collection Program 48
    2.1.1 Controlling the Interface 48
    2.1.2 Collecting Appropriate Data 49
    2.1.3 Advantages of Visual Basic 6.0 50
    2.1.4 Interface Development 50
        Developing the Visual Basic Application 51

The Development of Forms 51
2.1.5 The Code that accompanies the Interface 53
Displaying the Web pages 54
Developing the Data Files 55
2.2 Methods 57
2.2.1 Participants 57
2.2.2 Materials 58
The Hypertext Document 58
The Navigation Tools 58
The Tasks 64
The Information Search Task 64
The Free Recall Task 64
The Recognition Task 64
The Ratings Task 65
2.2.3 Procedures 66

CHAPTER 3 - PERFORMANCE IN HYPERSPACE

3.1. Issues related to Navigation Tools 68
3.2 Goals of the Performance Analysis 71
3.3 Results 71
3.3.1 Efficiency Measures 71
Effects on Speed of Navigation 71
Effects on Economy of Search 72
3.3.2 Effectiveness Measures 73
Effects on Free Recall 73
Effects on Recognition 74
Effects on Search Accuracy 75
3.4 Discussion 76
3.4.1 Summary and Implications 77

CHAPTER 4 - OVERT AND COVERT SPACES

4.1 The Behavioural Space 81
4.1.1 Goals of the Behavioural Analysis 81
4.1.2 Linearity of Use 82
4.1.3 A Qualitative Approach: Multidimensional Scaling 84
4.1.4 Results 86
Alphabetical 88
Hierarchical 91
Spatial 92
Spatial/hierarchical 93
Linearity 96
4.1.5 Discussion 96

4.2 The Representational Space 101
4.2.1 Goals of the Representational Analysis 103
4.2.2 Results...103
Instruction Conditions..103
The Relationship between Behaviour and Representation 104
4.2.3 Discussion 105
4.3 Summary and Implications 106
4.3.1 Performance and Representation 106
4.3.2 Behaviour and Representation 107

CHAPTER 5 - REPRESENTATION AND BEHAVIOUR – A SYNTHETIC APPROACH TO AN EARLY DEBATE

5.1 Introduction...109
5.2 Similarity Space Representations 110
5.3 Mental Imagery..111
5.4 Cognitive Mapping 113
5.5 A Synthetic Approach...116
5.5.1 Goals of the Study..118
5.5.2 Introduction to PDP networks...119
5.5.3 Value Unit Architecture...121
5.6 Simulation 1: Metric Spatial Judgments..122
5.6.1 Methods...122
Network Architecture..126
Input Unit Representations 126
Hidden Units...126
Output Unit Representations..126
Network Training...127
5.6.2 Results...128
Relating the map of Alberta to Hidden Unit Connection Weights 128
Relating connection weights to Hidden Unit MDS Spaces 131
Coarse Coding from the Hidden Unit Activations to Distance Ratings 134
5.6.3 Discussion..136
5.7 Simulation 2: Non-metric Spatial Judgments 137
5.7.1 Methods...138
Network Architecture and Training 139
5.7.2 Results...139
Relating the map of Alberta to Hidden Unit Connection Weights 140
Relating connection weights to Hidden Unit MDS Spaces 142
Coarse Coding from the Hidden Unit Activations to Distance Ratings 143

        Coarse Coding for Violations of Minimality 144
    5.7.3 Discussion 145
5.8 Summary and Implications 146

CHAPTER 6 – CONCLUSIONS

6.1 Conclusions 149
6.2 Impact to Practitioners 154
6.3 Limitations 155
6.4 Future Research 156

References 157
Appendix A 172
Appendix B 174
Appendix C 175
Appendix D 176
Appendix E 177

TABLES

Table 3.1. Means and Standard Deviations for the Average time across questions to complete the information search task for all groups........................72

Table 3.2. Means and Standard Deviations for the Average number of Pages accessed across questions to complete the information search task for all groups........................73

Table 3.3. Means and Standard Deviations for the Average number of Page Titles recalled for all groups........................74

Table 3.4. Means and Standard Deviations for the proportion of correct answers from the information search task for all groups........................75

Table 4.1 Correlations for average ratings between all groups........................103

Table 5.1. Distances between cities of Alberta, measured in kilometres...........123

Table 5.2.The distances from Table 6 converted into ratings on a 0 to 10 scale........................124

Table 5.3. Results of relating Alberta map distances between cities and hidden units the values of the connection weights feeding into the hidden units in Simulation 1........................130

Table 5.4. Results of relating connection weights to city distances from the MDS solutions obtained from the activity matrix for each hidden unit in Simulation 1. The table provides the maximum correlation, as well as the coordinates of the hidden unit in the space that produces this maximum correlation......................134

Table 5.5. Results of relating Alberta map distances between cities and hidden units the values of the connection weights feeding into the hidden units in Simulation 2 141

Table 5.6. Results of relating connection weights to city distances from the MDS solutions obtained from the activity matrix for each hidden unit in Simulation 2. The table provides the maximum correlation, as well as the coordinates of the hidden unit in the space that produces this maximum correlation 143

FIGURES

Figure 2.1. Administrative Login Form 53

Figure 2.2. Example of a Log Text File 56

Figure 2.3. Example of hypertext document page 59

Figure 2.4. a) Alphabetical List, b) Hierarchical list 60

Figure 2.5. Spatial Navigation Tool 61

Figure 2.6. Spatial/Hierarchical Navigation Tool 63

Figure 4.1. Three-dimensional solution for the Alphabetical Group – Axial Partitioning based on the divisions of the alphabetical list into thirds.....89

Figure 4.2. Alphabetical Facet Diagrams: (a) Dimension 1 and 2, (b) Dimension 1 and 3 90

Figure 4.3. Two-dimensional solution for the Hierarchical group axially partitioned by depth of hierarchy 92

Figure 4.4. Spatial two-dimensional solution: Axial partitioning and clustering according to three depth levels of hierarchy 93

Figure 4.5. Three dimensional MDS solution for the Spatial/hierarchical group – (a) Axial partitioning based on depth of hierarchy, (b) Radial partitioning by depth from Fungi center point 94

Figure 4.6. Three dimensional MDS solution for the Spatial/hierarchical group - Axial partitioning by breadth of the hierarchy 95

Figure 5.1. The locations of 13 Albertan cities taken from the MDS analysis of the ratings matrix of Table 5.2. The positions of the cities are very similar to their locations on a map of Alberta 125

Figure 5.2. A multilayer perceptron. The environment activates input units, which in turn produce activity in a layer of hidden units. The hidden units then produce a response by sending signals that produce activity in the output units. Links between layers are weighted connections; the values of the weights are determined by having the network learn through examples. Given two input city names, the network was trained to rate the distance between the cities on a map of Alberta 127

Figure 5.3. The locations of the six hidden units from the first simulation on the map of Alberta. These locations optimized the correlations between hidden unit connection weights and distances from the cities to the hidden units....131

Figure 5.4. The locations of the seven hidden units from the second simulation on the map of Alberta. These locations optimized the correlations between hidden unit connection weights and distances from the cities to the hidden units....142

CHAPTER 1: INTRODUCTION

For centuries, text-based literature has been the primary mode for disseminating and receiving information. The cognitive processes involved in reading and writing books have been studied extensively for decades as they have played a crucial role in the development of humankind. In the past two decades, hypertext systems (computer-based information systems of interconnected documents) have become increasingly pervasive in work and educational environments. With the advent of the World Wide Web, these information systems will heavily influence how we develop as a society in the future. Our individual competence with such systems will impact our everyday lives. Identifying the ways in which people conceptualize the computer environment (often referred to as "hyperspace"), interact with hypertext systems and access pertinent information is highly relevant to forecasting the full extent of that impact.

The study of hypertext systems is a multidisciplinary venture. Psychologists, computer scientists and educators all bring different approaches and goals to the effort. Psychologists wish to know how humans conceptualize hyperspace, specifically, how information in a hypertext format is perceived, retrieved and manipulated. Computer scientists want to know how to design an optimal interface for users. Educators are interested in how this technology can be applied to enhance learning. This can be advantageous as a broad knowledge base and diverse perspectives are brought to bear on the topic. However, it also results in the lack of a unified theoretical foundation for hypertext. Relative to many other areas, research on the cognitive aspects of using such systems is in its infancy; models and paradigms are still being developed. Towards the goal of understanding cognition in hypertext, researchers need to apply the same ingenuity, rigor and attention to detail in their methodologies as reading researchers have in the past to reveal the cognitive processes involved in various aspects of reading.

This thesis is a series of explorations regarding the cognitive processes that provide a foundation for the task of hypertext navigation. The overall goal of this thesis is to gain insight into the nature of the underlying mental representations that support hypertext navigation behaviours through the application of cognitive psychology and cognitive science methodologies.

1.1 A Roadmap

The current chapter is an extensive review of the literature on hypertext navigation from two fields of study: psychology and human-computer interaction. Part of this review focuses on the spatial metaphor; an early metaphor of hypertext use which compares hypertext navigation to aspects of human spatial processing in real-world, large-scale environments.

Chapter two describes the methods used to collect the human data for this project. It also includes a detailed description of the program designed by the author to collect a number of relevant types of data for the questions addressed in this thesis.

Chapter three presents the performance data collected; that is, both efficiency and effectiveness measures: speed, number of pages accessed, recall and accuracy. Chapter four examines the impact of different navigation tools on users' path patterns. This chapter makes use of multi-dimensional scaling for a qualitative examination of the general path sequences taken by each of the four navigation groups. These path patterns are compared to the configuration of the navigation tool for each group. This chapter also investigates if users' mental representations are impacted by the presence of spatial information in a navigation tool. Data from a ratings task is used to identify correlations between the users' perceptions of the document structure and their actual path patterns.

Chapter five begins with a theoretical discussion on the relationship between spatial mental representation and the behaviour generated by such representations. It then proceeds with two examples of spatial behaviour generated by artificial neural networks plus an in-depth analysis of the internal representations of the networks that guide spatial behaviours. An important question addressed in this chapter is: Must the underlying representations that produce spatial behaviours also be spatial in nature?

Chapter six provides a summary of the thesis findings accompanied by comments on the relation of this thesis to the spatial metaphor and broad theories of spatial cognition.

1.2 An introduction to hypertext

Generally, a hypertext system is an electronic system of interconnected units of information or text nodes. Many different types of hypertext systems have been designed, each consisting of its own specific structures that support the connections (links) between the user interface (the windows or pages a user views on the display terminal) and the underlying information system (database of text nodes). The user interface is sometimes referred to as the 'front-end' of a hypertext system and the underlying database as the 'back-end' of the system. In this paper, the term "window" will be used in reference to the screen image a user may be viewing and "node" in reference to the informational node that the window is connected to in the underlying database. The computer environment of hypertext documents has come to be referred to as "hyperspace". Hypertext documents that contain images, audio and video clips are referred to as "hypermedia" documents. In the hypermedia literature the terms "hypertext" and "hypermedia" are often used interchangeably (Bieber, Vitali, Ashman, Balasubramanian & Oinas-Kukkonen, 1997).

Each hypertext window contains icons that can be pictorial symbols or highlighted words (hotwords). A window can contain any number of these icons. The icon represents the link to the database for that particular piece of information. When clicked on with a pointing device (usually a mouse), the icon activates the link to the referenced node in the database and a new window is instantaneously opened containing the information in that node. Nodes that hold related information are linked. Some nodes will have a multitude of links, both to and from other nodes based on their relatedness, whereas other nodes may only have incoming links. It is important to stress that links are not merely connections between nodes but that they should specify the semantic relationships between nodes. The author of the hypertext document is the engineer of the links based on his/her semantic network, that is, his/her conceptions of the interconnectedness of ideas within the chosen material (Jonassen, 1986, 1989).

Links can be organized in two ways: in a hierarchical structure or in a network structure. In a hierarchical structure, nodes are only linked to the superordinate information above them or the subordinate information below them (Schneiderman & Kearsley, 1989; Jonassen, 1986). In a network structure, any node can be linked to any other. Multiple links can exist between superordinate and subordinate information (Schneiderman & Kearsley, 1989). Regardless of the specific structure, because of its general node/link relational structure, a hypertext system is non-linear (Conklin, 1987). That is, unlike the pages in a book, users can access different windows in different orders. There is no predetermined sequential organization of the underlying nodes. Typically, a hypertext system will also have a backtrack function to allow the user to access previously viewed windows (Nielsen, 1990).

To summarize, the defining features of hypertext are: a large collection of informational nodes connected via machine-supported links, which support massive cross-referencing and easy access. The term "hypertext" will be used as a general term for systems which entail these features regardless of their individual differences.

1.2.1 Goals in hypertext use

Different hypertext users have different goals, which produce different navigational strategies. Users may be searching out a very specific piece of information that they have already accessed, looking for a specific piece of information that they do not know how to find or they may be exploring the system with no specific informational goal. Conklin (1987) outlines the three methods of conducting an informational search in a hypertext system; (1) the user can visit a set of related windows or pages through the traversal of links, (2) the user can navigate through the document using a graphical browser which provides a graphical representation of the network structure, or (3) the user can activate a queried search where the system searches all documents to find any instance of a keyword or key phrase that is entered. The first two methods described above are relevant to documents that are not overly large and that can be viewed exhaustively in a reasonable time period. The query method is necessary for documents that cannot be exhaustively examined because they contain too many nodes. This thesis focuses on the first two methods of hypertext use, and therefore on interaction with smaller documents, rather than the queried search method.

1.2.2 Advantages and Disadvantages of Hypertext

The main advantage of hypertext over traditional text is that it is a high-speed method of accessing vast quantities of information. As mentioned earlier, there is no predetermined sequential organization of the underlying nodes. This provides an increased flexibility in the ways that users can experience the information in the document compared to traditional text-based presentations. The user is free to make link connections in any order based on their own focus and interpretations of the material, toward their own end goal. However, this same feature that provides such individual freedom for the user also comes with costs. For example, at any point in time, the user only sees the current window and the links represented by icons or highlighted text, leading from that window to other related windows. This means the user must remember the windows already viewed and hypothesize about the overall size, structure and complexity of the entire hypertext network, as it can never be viewed in its entirety.

Because of this, there are two related categories of problems associated with using hypertext systems: cognitive overhead and disorientation.

Cognitive Overhead

Unlike traditional text, hypertext does not have conventional cues (e.g., chapters, sections) that lead the reader through a document (Gygi, 1990). With traditional text documents, format often indicates the type and purpose of a document and gives the reader hints about where important information will be found (e.g., a legal contract is easily distinguished from a grocery list). In hypertext, no such conventions have been established. In addition, there are multiple routes or paths to access a specific piece of information in the network. This increased flexibility results in users being responsible for decisions about which information to access (which links to select) and in what order. To this end, the cognitive resources of the user must be allocated toward several goals concurrently. The term ”cognitive overhead” has been used to refer to the amount of cognitive resources necessary to successfully complete an informational task in hypertext (Conklin, 1987).

With hypertext, users must evaluate information currently viewed, must recall information previously viewed as well as its location relative to other viewings, and must create a plan to navigate toward desirable information (not only deciding which windows to access but also which tools to use to execute the plan). Kim and Hirtle (1995) label these cognitive tasks as: (1) navigational tasks: planning and executing routes through the network; (2) informational tasks: reading and understanding the contents presented in the nodes and their relationships, for summary and analysis; and (3) task management: co-coordinating informational and navigational tasks. The cognitive demands that these tasks require are compounded by increasingly complex and voluminous hypertext structures (Schneiderman & Kearsley, 1989; Conklin, 1987; Jonassen, 1989; Tsai, 1989; Ransom, Wu & Schmidt, 1997). Degraded performance occurs when a hypertext user's cognitive resources are overwhelmed by these demands.

Disorientation

One of the symptoms of cognitive overhead is disorientation (Conklin, 1987; McLeese, 1989; Ransom, Wu & Schmidt, 1997). It is well documented that hypertext users will often report "feeling lost" when performing poorly on an informational search. (Gygi, 1990; McLeese, 1989; Edwards & Hardman, 1989; Jonassen, 1988). Based on subjective reports, Edwards and Hardman (1989) characterized this phenomenon via three categories of the users' experience: 1) the user does not know where to go next, 2) the user knows where to go but not how to get there, 3) the user does not know where they are in relation to the overall structure of the document. Based on empirical measures, Foss (1989) proposed three types of disorientation problems and their symptoms: 1) Navigational disorientation problems, caused by a lack of knowledge of organization and extent of the document as well as unfamiliarity with navigational tools. The symptoms of navigational disorientation are looping, inefficient pathways and query failures. 2) The embedded digression problem caused by high cognitive demands that lead to difficulties in planning, managing and executing digressions. The symptoms for embedded digression are a disorganized screen layout, having many windows open concurrently and engaging in excessive backtracking. 3) The art museum problem caused by the high cognitive demand of viewing many windows superficially without attention to detail or relatedness. The symptoms of the art museum problem are short reading times and restricted search paths.

The problems of disorientation and cognitive overhead were identified by early hypertext researchers, (McLeese, 1989; Edwards and Hardman, 1989; Jonassen, 1988; Schneiderman & Kearsley, 1989) but continue to be unresolved (Ransom, Wu & Schmidt, 1997; Dias & Souza, 1997). It has been suggested that unless disorientation and cognitive overhead can be effectively reduced the ultimate usefulness of hypertext will be severely limited (Conklin, 1987; McLeese, 1989; Ransom, Wu & Schmidt, 1997)

1.3 The Spatial Metaphor

Given the apparent difficulties that arise in hypertext use, an initial question to address in investigating the cognitive processes involved is - what, essentially, is the psychological nature of the hypertext environment? The identification and labeling of the two problems of disorientation and cognitive overhead are a direct result of how we have come to conceive of the environment that hypertext creates. Disorientation is a cognitive phenomenon that occurs in an environment where it is necessary to maintain a sense of one's own location relative to the location of other objects and to maintain the relative location of these objects to each other. Cognitive overhead in hypertext is the amount of cognitive resources that are required to successfully interact with the type and number of structures within that environment (the type of structures dictated by our conception of the environment). Thus these terms reflect the notion that hypertext exists in a spatial environment, an idea that is central to an early metaphor for hypertext use, the spatial metaphor.

The spatial metaphor arose out of a need for a common language to converse about the hypertext environment and, subsequent to that, out of the need for a framework for designing tools and hypertext features that create optimal navigating environments. Its basic premise is that locating information in hyperspace has similar psychological features to navigating in real-world space. Canter, Rivers and Storrs (1985) have been attributed with the emergence of the metaphor (Dias & Sousa, 1997), which has been advocated by a number of researchers (McDonald & Stevens, 1998, 1999, Hammond & Allinson, 1987; Kim & Hirtle, 1995; Leventhal, Teasley, Instone, Rohlman & Farhat, 1993). However, other authors have expressed concern over the rapid and seemingly overwhelming acceptance of the metaphor (Shum, 1990; Jones & Dumais, 1986; Stanton & Baber, 1994; Dias & Souza, 1997; Mayes, Kibby & Anderson, 1990). Navigational tools modeled after real-world navigational aids appear to enhance performance on some measures (Kim, 1999; McDonald & Stevenson, 1998,1999; Beasley & Waugh, 1995; Schroeder & Grabowski, 1995) but not on others (Stanton, Taylor & Tweedie, 1992; Dias & Sousa, 1997; Wenger & Payne, 1994; Leventhal, Teasley, Instone, Rohlman & Farhat, 1993; Reynolds & Dansereau, 1990). Correlational data on the relationship between spatial ability and performance has shown increases in speed of performance with higher spatial ability (Campagnoni & Ehrlich, 1989; Vicente & Williges, 1988) but marginal effects on the accuracy of performance (Leidig, 1992). Therefore, the empirical evidence to date has not resolved the issue of how or to what degree spatial processing is involved in the efficiency and effectiveness of hypertext use.

The spatial metaphor seems easily understood by novice users. However, if it is not an accurate metaphor for the mental processes that humans engage in as they interact with hypertext or if we do not understand the aspects of the comparison that make the metaphor beneficial, then as more and more complex, powerful systems are developed this metaphor will eventually be less useful. Is the spatial metaphor intended as a theory-constitutive metaphor, that is, as a psychological theory of how hypertext navigation occurs or does the spatial metaphor only function as a tool for interface design?

Having presented an overview of hypertext systems, and a brief description of how users interact with them, it should be clear how firmly entrenched the spatial metaphor is in the hypertext literature. The use of terms related to space is solidly embedded in the topic of hypertext. Terms such as 'hyperspace', 'environment' and 'navigation' are difficult to avoid as, at this point in time, they provide for the user the most direct access to meaning at an introductory level. To this end, the spatial metaphor has been useful and may continue to be useful. However, beyond the introductory level, is the spatial metaphor impeding the development of a more cognitively accurate model, one that might ultimately provide more benefit to the user?"

Before addressing this larger issue, it is prudent to attempt to make explicit the ways in which the terms "spatial" and "metaphor" are used (or not used) in relation to hypertext.

1.3.1 The Role of Metaphor

The research on metaphor and its related term, analogy, is vast and complicated. The following section is a brief review of perceptions of metaphor in both the psychological and human-computer interaction literature.

In the psychological literature, Lakoff (1994) describes metaphor as a mode of thought that can be understood as a tightly structured mapping from the source domain to the target domain. The source domain is a conceptual schema that we have constructed based on our experiences in a particular environment. In the case of the spatial metaphor, the source domain is physical space and the target domain is hyperspace.

Holyoak (1996) considers metaphor, like analogy, a method toward understanding a new domain in terms of a source domain, but metaphor involves connections that move beyond surface category structure. " Many metaphors are based on deeper relational and system mappings." (p. 217).

Holyoak stresses the importance of the evaluation of these mappings based on similarity, structure and purpose between the source and target domains.

According to Pylyshyn, the role of a metaphor for scientific explanation is not to fully explain all aspects of a complex environment but to provide a framework on which to reference new, vague and disconnected ideas about that environment, phenomenon or concept. This framework is a means toward the development of a more thorough understanding of something new. Pylyshyn (1993) describes the acquisition of new knowledge using Piaget's terms of accommodation and assimilation. When we first encounter new information, we attend to aspects that are most similar to the schemas we've already constructed and categorize the new information as belonging to the most similar schema (assimilation). If the new information is so dissimilar that it cannot be mapped onto an existing schema, we then adjust and expand the structure of the schema to include this new knowledge (accommodation). However, Pylyshyn maintains that it is important to provide some degree of explanation via the metaphor (a powerful metaphor) rather than merely description, which may leave the user with a groundless impression of having been provided with explanation (an impotent metaphor). The value of a metaphor decreases if the user, upon closer inspection, realizes that the metaphor has few shared properties with the new situation the person is attempting to understand. Surface similarities may prove of limited value as functional and conceptual dissimilarities are uncovered.

In hypertext, as the user becomes more proficient the comparison between hyperspace and physical space may break down. The spatial metaphor may be useful as a pedagogical tool to introduce the novice user to the system and to aid in the initial organization of his/her thoughts about this new environment but may also impede the user from developing a more in-depth and explicit conceptualization of the hypertext environment. As Pylyshyn notes (1993),

> "One must distinguish between the general programmatic enterprise of trying to illuminate a new phenomenon and the much more demanding goal of establishing the validity of an explanatory theoretical principle" (p. 548).

Therefore, one goal in metaphor use within psychology is to define and explain the relations between the source and the target domain in the metaphor, and the cognitive processes involved in executing the task associated with that metaphor. It seems the spatial metaphor has been prematurely accepted as a theoretically plausible metaphor for the cognitive processes involved in hypertext interaction.

Unlike psychological theory, in the human-computer interaction (HCI) literature the end goal is to provide a metaphor for the user that helps them understand the purpose and function of the facilities of a specific application rather than to provide the user with an explanation of the global environment of hypertext. Metaphors can be applied to an application with a narrow purpose (e.g., the typewriter metaphor for word-processing) or a broader purpose (e.g., the general desktop metaphor used in many types of applications). In the HCI literature, some authors are careful to explicitly present the spatial metaphor as a design tool for a specific application. An example of this is Hammond and Allinson's (1987) travel holiday metaphor, which includes choices for go-it-alone travel or guided tour travel for a specific hypertext system. Other authors present the spatial metaphor as if they accept the assumption that it is a psychologically plausible account of how people think about hyperspace in general (Kim & Hirtle, 1996; McDonald & Stevenson, 1997).

Interface designers expect and accept the presence of metaphorical mismatches. Carroll, Mack and Kellogg (1988) emphasize the usefulness of mismatches as does Hammond & Allinson (1987). Hammond and Allinson contend that a metaphor need not be accurate at all levels of description to be useful,

> "Furthermore, a metaphor which is self-evidently inappropriate at some levels of description will be more successful than one which is ambiguous: no attempts will be made to invoke mappings at the inappropriate level" (p. 83).

An example of a metaphor mismatch is exhibited in the metaphor of word processor as typewriter. On a typewriter, the space bar moves the typing position whereas on a computer keyboard the space bar inserts a space. Novice users are often confused when the spacebar does not produce the same result on a computer as it does on a typewriter (Hammond & Allinson, 1987).

Although mismatches are expected, these researchers also advocate a more systematic assessment of metaphor to interface design. Carroll et al. (1988) argued that three aspects of analysis must be undertaken to properly assess the value of an interface metaphor: operational, structural and pragmatic analyses. Operational analysis entails the measurement of behavioural effects (e.g., improvements in performance and learning). Structural analysis involves making explicit the mappings of referent domain to target domains, a description of primitives and relations between primitives in the referent and target domains. Structural analysis is necessary to determine the scope of the metaphor. Pragmatic analysis examines the practical context of the metaphor. This includes examining the user's goal associated with using the metaphor as well as the impact on the user of the flaws in the metaphor (e.g., mismatches). This is similar to Holyoak's focus on similarity, structure and purpose in a psychological approach to metaphor validation. Carroll et al. emphasized that all three types of analyses are necessary to fully understand the dynamics between metaphor and user and, hence, to assess the overall usefulness of the metaphor.

> "An adequate analysis of interface metaphors needs to rest on an empirical task analysis of what users actually do, not an abstract normative analysis of what they might do"(p. 73).

Psychological and HCI researchers espouse similar types of metaphor assessment but towards a different focus; respectively, understanding underlying information processing mechanisms vs. usability. Both the psychological and HCI approaches seem plausible for the evaluation of a broad metaphor like the spatial metaphor. Nonetheless, in a survey of the literature, it seems neither approach has been used in a systematic way to closely examine the hypertext/real-space comparison. The mappings between hypertext features and real-world features have not been empirically specified nor have particular cognitive processes involved in the use of hypertext (e.g., degree of spatial processing) been identified in detail.

1.3.2 Mental Models vs. Mental Representations

At this juncture, it is relevant to examine another disparity between the hypertext literature and psychological literature. The term mental model is used fairly extensively in both areas. For human-computer interaction researchers, notions of mental models encompass a multifaceted representation (Sasse, 1992). This includes mental representations not only of the content space but, more pointedly, models of how the system works, what actions must be executed sequentially to complete a task, what is the purpose and functionality of screen elements. The goal associated with this definition of the user mental model is to learn how the system functions, the actual content not being as relevant. This is very different from the goal of retrieving, understanding and forming mental representations of the information that comprises the content space of the document. Mental models for educational and psychological purposes tend toward this latter description. To keep this distinction in mind, I will use the term mental representation rather than mental model.

In addition, psychological researchers attempt to distinguish clearly how mental representations impact behaviour at three specific levels: behaviour originating with the user, the behaviour elicited by the environment and behaviour generated by the interaction of the two. Therefore, to fully understand a phenomenon, the properties of the behaviour, the properties of the environment and the properties of the mental representation of that environment must be individually specified. The underlying mental representations that users develop of the content of documents may be spatial. This does not necessarily dictate that users' behaviours are spatially organized. This distinction between representation and behaviour is often overlooked in HCI studies.

1.3.3 Relating Space to Hypertext

While investigating the use of a spatial metaphor, it is also important to consider definitions of space itself. For example, what are spatial properties and how do these properties map onto "hyperspace"?

Metric Models of Space

Real-world, everyday, physical space is generally described in terms of geometric properties, specifically Euclidean or metric properties. Three-dimensional space is denoted by the geometric relationships between three points, X,Y, and Z. These relationships include such properties as direction and distance. Metric space adheres to three geometric axioms: minimality, symmetry and triangle inequality (Blumenthal, 1953). These axioms will be described in detail presently.

Stanton and Baber (1994) express concern that this everyday definition of space has been inappropriately applied to hyperspace (which they refer to as 'electronic space'). They state, " the basic concept of hypertext in computer science terms, is an n-dimensional space which can be traversed by moving through links. 'Space' as used in this context has a well-defined meaning, as the collection of objects and activities contained within a specific domain, similar to its use in finite state architecture. However, much current research in hypertext appears to use the term 'space' in its everyday sense, that is, as a physical relationship between objects" (p.235-236).

In the case of hypertext, it appears that physical distance is being equated to some degree with psychological distance. What exactly is psychological distance? In hyperspace, one way of conceptualizing psychological distance would be an estimate of the degree of semantic relatedness or similarity between two windows.

Geometric models, like the Euclidean model of physical space, are often used in the psychological literature to describe psychological processes. Beginning in the early 1960's, many psychological researchers (e.g., Shepard, 1962; Kruskal, 1964; Guttman, 1971; Carroll & Wish, 1974) asserted that the similarity of objects could be modeled using geometric properties. More explicitly, the similarity of two objects could be represented as points in some coordinate space and that a metric distance function could be used to compute the dissimilarity between two points. A metric distance function assigns a distance to each pair of points based on adherence to the same three axioms that define physical space: Minimality, symmetry and triangle inequality. The minimality axiom implies that the similarity between an object and itself is the same for all objects. The assumption of the symmetry axiom is that the distance between two points is the same no matter the order, that is, the distance between A and B is the same as the distance between B and A. The triangle inequality axiom assumes that the shortest distance between two points is a straight line. That is, if A, B, C are any three points in space, then the sum of any two distances (e.g., AB, BC, AC) is greater than or equal to the third distance. Applied to similarity theory, it implies that if A is similar to B, and B is similar to C, then A and C cannot be very dissimilar from each other (Tversky, 1977).

Contrary to the above view, Tversky (1977, 1982) argued that although similarity relations may be represented using geometric models including the axioms presented above, these representations should not be viewed as psychological theory. Tversky contended that the minimality axiom is not fulfilled by some types of similarity measures. For example, the probability of judging two identical stimuli as “same” rather than “different” is not constant for all stimuli. In experiments on recognition of simple objects, sometimes an object is identified as another object more frequently than it is identified as itself. According to Tversky, the symmetry axiom does not hold for similarity relations either. Similarity judgments can be regarded as extensions of similarity statements, that is, statements of the form “a is like b”. In such a statement, a subject and a referent are evident. The statement is directional as it begins with the subject and ends with the referent. To reverse the statement is to lose its meaning; "b is like a " is not the same as "a is like b".

As for the triangle inequality axiom, Tversky stated that the implication mentioned above (if A is similar to B and B is similar to C, then A and C cannot be very dissimilar from each other) sets limits on the similarity between A and C, as the similarity between A and B, and B and C, constrains the similarity between A and C. Tversky provided an example that challenges the psychological validity of the triangle inequality assumption . He makes comparisons between the following pairs of countries, Jamaica - Cuba, Cuba - Russia, Jamaica - Russia. The dimension that defines similarity between Jamaica and Cuba (geographic similarity) is not the same as the dimension that defines the similarity between Cuba and Russia (political similarity, at the time). Therefore, assigning a degree of similarity between Jamaica and Russia is uninformative, as comparisons are not made on dimensions that are valid between the two countries.

How well do geometric properties map onto hyperspace? In hyperspace, distance could be thought of as the number of steps between windows (e.g., number of mouse clicks), the number of layers to traverse in a hierarchical system or the amount of time it takes to access one window from another (Shum, 1990). However, it could be said that these variables are only epiphenomena of the underlying semantic relationships of the material as perceived by the author, the person who has created the structure of the hypertext document. The links between windows are constructed by the author based on his/her perceptions of the relatedness of meaning between the content of each window. In this sense, the links are somewhat arbitrary in that the degree of relatedness perceived by the author may be very different from the degree of relatedness perceived by the user. Also, depending on the order in which the user has accessed certain windows, these variables may not consistently represent, from user to user, the underlying semantic relationships between windows.

The question therefore becomes, in the hypertext context, how does meaning map onto space? One implication of the spatial metaphor is that the degree of semantic similarity between windows can be psychologically represented as a distance. Windows that are similar in meaning are "close together" in the "information space", windows that are not are "far" apart. This, in turn, implies that semantic similarity can be described using geometric properties just like Euclidean geometric properties define real-world space, hence, connections between windows can be conceptualized as "space". Although a 'metric' distance is not applied to each pair of hypertext windows, because of the comparison to real space prompted by the spatial metaphor, there seems to be an underlying assumption that some sort of distance relations exist.

However, as outlined above, Tversky provides some compelling reasons for why psychological estimates of similarity should not be represented geometrically, that is, as a metric distance, for all stimuli in all situations. The hyperspace environment may violate the axioms necessary for a geometric similarity model. Previous psychological studies illustrate that the minimality axiom is not always accurately adhered to with simple stimuli. Referring to semantic similarity between hypertext windows, this problem is intensified due to the complexity of each window and the sheer number of windows being accessed in large documents. The symmetry axiom may also violated. Once a new window is accessed, the information the user encodes there goes with him/her back to the previous window, where that window is likely to be viewed in a slightly different manner. Therefore, knowledge-wise, the initial “distance” from the first window to the second or subsequent windows is not the same as the ”distance “ from the second window back to the first. The ‘distance’ has been altered merely through the accessing of the second window and its contents. As for the triangle inequality axiom, different windows may or may not share the dimensions that other pairs of windows share just as, in Tversky’s example, Jamaica, Cuba, and Russia do not share the same dimensions.

This discussion of similarity theory has been applied toward the validity of using a geometric model for semantic similarity in hypertext just as a geometric model is used to describe real-space (Euclidean properties). The depth of comparison in this example may be excessive, but it is a demonstration of how unspecified the parameters of the spatial metaphor remain. We are still left with an ill-defined and vague sense of how space is related to the hypertext environment and the elements within it.

Spatial Primitives

A less formal way of defining space is through the decomposition of its characteristics into a set of basic primitives. Primitives can be thought of as both elements and relationships that are necessary to convey the fundamental attributes of a phenomenon. Golledge (1994) identifies and describes four spatial primitives: 1) identity, 2) location 3) magnitude and 4) time. Golledge describes providing identity ”as the process of equating an occurrence with a name or label. The purpose behind this differentiation is to allow occurrences to be recognized and their uniqueness evaluated”

(p. 104, 1994). As Golledge notes, a problem with identity is that, at times, two labels may be given to the same entity (e.g., a street is known by different names to different people). In hypertext, if an identity is not universally acknowledged, not only may different users be unable to recognize an item as the same item but also individual users may confuse different items.

The second primitive, location, is denoted as "information about where an occurrence exists within the totality of the environment" (p.104, 1994). The way location is specified can be different depending on the environment being considered. The spatial terms and reference points that describes an entity's location will be different if that entity is a continent rather than a country, a city rather than a building or an object. Locational information includes information that describes relations between an entity and other entities within that specific environment. Understanding the idea of location in hypertext has several obstacles. How do we describe different categories of environments (e.g., what is global, what is local?) How do we specify location within that environment? An item in hyperspace has no absolute location or even a consistent relative location.

From the property of location, several derived concepts emerge: 1) distance, 2) direction, 3) sequence and order, and 4) connection and linkage. Golledge defines distance as the " interval between the locations of occurrences" (p.106, 1994). To quantify distance, location must be clearly specified. Direction can only be specified within an established frame of reference. Sequence and order are dependent on definitions of distance and direction as well. Connection and linkage are related to principles of proximity and similarity, which again, are inextricably tied to the specification of location. The relation between hyperspace and these derived concepts cannot be determined until the relation between hyperspace and location is clearly understood.

The third primitive, magnitude, refers to the size of an identity. In instances where we can define " how much" with a single number the measurement of magnitude is straightforward (e.g., how many cows are in a field). In other instances, Golledge concedes that magnitude is difficult to measure. For example, do we define a city by the area covered, the size of the population or by the boundaries of specific activities? With hypertext, do we measure magnitude by the number of nodes in a document, the degree of linkage within the document or the degree of linkage to other documents within a network? It is not clear how magnitude may be defined in hyperspace.

The fourth primitive, time, denotes the permanency of an occurrence. Does an occurrence exist at a later point in time just as it exists in this moment? Does an occurrence retain its location over time? The relation between time and location is particularly tenuous in hypertext as we may not even be able to denote location itself. As well, large systems of hypertext documents such as the Web are dynamic. Using a particular site as a landmark for locating other sites will be ineffective as sites can be removed, linking can be changed and the general appearance of the site may be altered. Beyond the issue of location, unlike real-space, navigation time in hypertext is not correlated to "distance". Changing locations between windows that are unrelated can be achieved just as quickly as moving between windows that are closely related.

Whether we approach the hypertext/real-space comparison from the axioms of Euclidean space or a list of spatial primitives, identifying discrepancies between the two environments will inform us about the value of the spatial metaphor towards enhanced information access.

1.4 Spatial Cognition

If spatial cognition plays an important role in how people interact with hypertext, comparisons between behaviours in hyperspace and behaviours in physical space should be possible on several levels: Do people hold spatial representations in memory of the hypertext environment? Do novice hypertext users go through the same process that people in novel physical environments do? Do children show similar patterns of development in hypertext navigation as real world navigation? To address these questions we must have a brief look at spatial cognition and how it develops in real-world environments.

The study of spatial cognition is characterized by diverse approaches in the choice of settings and paradigms. For example, subjects have been asked to find their way (e.g., Passini, 1984), to locate an object in real space (e.g., Cornell & Heth, 1983) as well as to imagine how an object might look if it was rotated to a different orientation (e.g., Shepard & Metzler, 1971). Small-scale environments such as a single room where all contents are visible have been studied (e.g., Acredolo, 1978) and, conversely, large environments, such as the layout of a campus or city, where all locations are not within view (e.g., Cornell, Heth & Alberts, 1994). Researchers have been interested in people's knowledge of particular places as well as their knowledge of abstract spatial concepts.

Subsequently, there are a variety of definitions of spatial cognition, each prompted by and constrained by the environment under investigation.

In general terms, spatial cognition is the knowledge and mental representation of the structure of space and the location and relation of objects or entities within that space, along with the knowledge of when and how to use this stored information and the ability to think about the mental representations themselves (Liben, 1981). This definition includes two parts: the informational system itself and the individual's ability to access and make use of that information. The informational system has been referred to as a cognitive map, a term first used by Tolman (1948). In this paper, the term cognitive map will be used to refer to the informational system that is part of and supports spatial cognition.

1.4.1 Cognitive Maps

Cognitive maps are long-term memory representations of spatial information. As we cannot view cognitive maps directly, we must infer their features from behaviour during specific tasks. Consequently, empirical evidence involves these representations as they are applied to the performance of a particular action. External forms produced by behaviour (e.g., sketch maps, verbal protocols, quantitative and structural analyses), are often called spatial products, a term originally used by Liben (1981). Spatial products give us some indication of the nature of underlying spatial representations. Behavioural evidence suggests what the rules and operations are that govern the use of spatial representations (Presson & Somerville, 1985).

The term cognitive map often elicits the assumption that this representation is an image-based, map-like mental representation. The actual form of cognitive maps has been highly debated. Some researchers have suggested that the cognitive map is analogous to a cartographic map, a Euclidean model of the world (Lieblich & Arbib, 1982). This is not to say that cognitive maps are exact replications of features on a map. They have been described as "schematic, sketchy, incomplete, distorted, and otherwise simplified and idiosyncratic" (Kaplan, 1973, p. 276).

Other researchers have suggested a conceptual-propositional, rather than image-based, theory of spatial information encoding (for review see Golledge, 1997). In this view, both verbal and visual information are stored as abstract conceptual propositions. Concepts are embedded in a propositional network and conceptual meanings are partially dependent on their relations to other concepts in the network. If concepts are close together in the propositional network, they will function as cues for the retrieval of one another. Alternatively, some researchers (Kitchin, 1994; Golledge, 1997) have argued that the term "map" should only be thought of as a hypothetical construct for the spatial system that we accept as existing and influencing our spatial decisions.

Aside from the debate about the form of cognitive maps, Kitchin (1994) summarizes the types of information that cognitive maps contain. Cognitive maps are selective, subjective representations of the environment that preserve the locational and attributional features that are most informative for the individual perceiving that environment. The attributes that designate the structure of the map (landmarks or labels) are markers for action that should take place at that point in space and time. According to

Kitchin, this is an important aspect of the cognitive map, that it is constructed to combine the concepts of place and action. Cognitive maps are individual constructions for individual events, there is no singular cognitive map that retains all the spatial information and knowledge we've acquired. They are dynamic, that is, our knowledge (cognitive map) influences the experiences we have and our experiences alter our knowledge (cognitive map).

More recent research into the nature of cognitive maps has suggested that their hierarchical structure is important. Hierarchies are organized by semantic and conceptual relatedness not by spatial relations. Several studies provide evidence that cognitive maps have a hierarchical structure, that is, entities are encoded at different levels or regions (e.g., particular objects in relation to sets of objects) based on their conceptual similarity or dissimilarity rather than their actual spatial location.

For example, Stevens and Coupe (1978) found that people are influenced by superordinate spatial relations of geographical locations. They asked subjects to judge the spatial relations of different cities (e.g., Is Montreal north or south of Seattle?). People often misjudged the location of different cities based on the state or country the city is situated in (e.g., Montreal is in Canada and Canada is north of the United States, therefore Seattle must be south of Montreal). The superordinate category of country influenced judgments about the subordinate category of cities, irrespective of their spatial locations.

McNamara, Hardy and Hirtle (1989) submitted subjects' recall protocols of a spatial layout to an algorithm that constructs hierarchical trees consistent with the internal organization of the protocols. They found that subjects' distance estimations and spatial priming was dependent on whether pairs of objects appeared in the same subtree or different subtrees of the hierarchy rather than each objects spatial location in the total array.

Both of the above studies highlight the importance of categorization and meaning in the organization of mental representations of spatial information. If hierarchical structure is important to the retrieval of spatial memories of objects in physical space, it is plausible that hierarchical structure is even more relevant than spatial position to the recall of objects in hyperspace where objects have no absolute position or even consistent relative position.

1.4.2 Spatial Knowledge Acquisition

In a highly influential article, Siegel and White (1975) suggested that the ability to act efficiently within large-scale environments is based on the mental integration of multiple views. Because large environments cannot be viewed all at once, some type of mental spatial representation must be constructed in order to comprehend the entire space. Siegel and White proposed that the acquisition of these spatial representations follows a distinctive course, as a series of stages, landmark recognition being the first in which individuals learn to discern and remember separate landmarks. In the second stage, landmarks are coordinated in a sequence, which results in path or route representations. This route knowledge is comprised of both spatial and temporal relations between landmarks. As route representations become more refined, a representation of the entire space forms, a survey representation, which includes the relations between landmarks and routes. Thus, with increased familiarity of the environment, spatial representations gradually incorporate more sophisticated versions of landmark, then route, then survey knowledge; survey knowledge being qualitatively different from the first two because it incorporates the understanding of metric distance relations between landmarks and routes. Siegel and White's view remained the dominant framework throughout the 1970's and 1980's and has been cited in the hypertext literature (Dillon, McKnight & Simpson, 1990; Edwards & Hardman, 1989; Leventhal, Teasley, Instone, Rohlman & Farhat, 1993)

More recently, Montello (1993) argued that the final acquisition of survey knowledge is a quantitative rather than a qualitative change. He asserts that metric knowledge is being encoded right from the earliest learning stages and that survey knowledge merely appears as a qualitative outcome of the integration of spatial knowledge, including metric knowledge, about separately learned places. Montello describes spatial knowledge as consisting of both nonmetric and metric information beginning with the first exposure to a new place. He also asserts that some types of spatial information are stored in a linguistic format and comes to the conclusion that " there exist two or more distinct subsystems that represent spatial information in different formats"(p. 17, 1993). Montello cites several authors who present converging evidence and arrive at comparable conclusions. For example, McNamara, Halpin & Hardy (1992) referred to three formats of spatial knowledge – temporal strings, propositions and metric spatial representations. Kuipers (1982) proposed that spatial knowledge is stored in many disconnected components that include separate metric and topological components.

Kosslyn (1987) suggested that the right hemisphere of the brain is specialized for processing coordinate spatial representations while the left hemisphere is specialized for processing categorical representations.

Three questions emerge from Montello's position: a) If metric knowledge is so pervasive throughout the acquisition of knowledge of large-spaces, how relevant is real world navigation to hypertext as metric distance may not be a property of hypertext?, b) Do spatial aids prompt the user to develop some form of mental representation that includes metric distances between elements? c) Integration of multiple types of spatial information is cognitively demanding (Golledge, Ruggles, Pellegrino & Gale 1993), therefore, what is gained by adding components to hypertext (e.g., spatial maps) that may force another knowledge subsystem (in this case metric knowledge) to be engaged?

An alternative perspective is that navigation may be route-based, that is, route knowledge is not a prerequisite for or subset of survey knowledge but rather route and survey knowledge are acquired and represented separately (Hirtle & Hudson, 1991). Hirtle and Hudson found that subjects could acquire route knowledge from both slide presentations and maps but could only acquire survey knowledge from maps. This suggests that, with hypertext, even without global survey knowledge of a document, users should be able to encode sequences of landmarks into routes and integrate separate routes for a local representation of parts of the document. However, even with a route-based explanation of hypertext use, we still encounter the problem of distance definition. A route is defined by the sequence of and distances between landmarks. In hypertext, distance is not easily measured and is inconsistent from one traversal of the document to the next.

1.4.3 A Preliminary Comparison of Hyperspace and Physical Space Features

Metric distance may be one property that is not shared by hyperspace and physical space. In the following section several other possible disparities are suggested as examples of the types of issues that need to be investigated to evaluate the appropriateness of the spatial metaphor.

Perceptual Cues

As mentioned previously, landmarks are important in navigation because they are a cue for a specific action to be carried out. In real-space, the surrounding environment provides additional retrieval cues for that action. In hypertext, we view only one window at a time. If the window itself is not a strong enough cue to retrieve the action from memory (e.g., select a particular icon to another window) and there are no navigational aids to help orient the user, the user must make a random selection of icons.

In contrast, real-world environments have an abundance of perceptual information, that is, they provide a rich collection of features and attributes on which we can base spatial decisions. Perceptual information allows us to make predictions about the future positions of objects in the environment as we move through space. For example, in a city setting, distal landmarks such as tall buildings are occluded by closer buildings. As we move through space, the degree of occlusion changes giving us clues as to the future position of the buildings that are not entirely in our view in that moment.

Hypertext environments are featurally impoverished compared to real-world environments. This may make the attributes that are present more salient but it also provides fewer opportunities for the user to establish multiple reference points that designate the users position along a route. There is no perceptual information that is inherent to the hyperspace environment that helps us make predictions about the future positions of objects.

The Nature of Landmarks

As navigational elements, in real-space, landmarks have no inherent meaning. Their relevance is learned and defined by the actions attached to them. Landmarks have a sequential relationship but not a semantic relationship. For example, a tree has no shared meaning with a car unless an action pairs the two together such as, " I must walk to the car and turn left towards the tree". In hypertext, landmarks, such as individual windows, have their own semantic content that is related to other landmarks aside from their possible sequential or spatial relationship.

A sequence of landmarks makes up a route. In real-space, the length of that route and the spaces between landmarks can be defined by a distance. Distance in real-space is correlated with time, that is, the longer the distance the more time it takes to traverse it. In hyperspace , this correlation does not hold, partly because we have no way of accurately measuring distance but also because hypertext is dynamic. Distances do not remain static but change depending on the order of access the individual uses at any given time or the different preferences expressed by different users.

Spatial Mental Representations

Survey knowledge in a large hypertext document can only be achieved through navigational aids such as maps and content lists. This means that it is possible that survey knowledge is mentally represented in a different form than route knowledge where the user has visually experienced the pages. If survey knowledge is encoded in a different form than route knowledge, it must undergo some kind of mental manipulation so as to track it to route knowledge. This increases the cognitive demands of the task.

Contrary to this, in real world navigation distal landmarks and the general lay of the landscape are visible. Through these cues, survey knowledge can be partially acquired in the same manner as route knowledge; it is an expansion of the cognitive map developed for route knowledge. The accumulation of metric and nonmetric information occurs from the outset, that is, survey knowledge begins to be encoded in the same manner, at the same time as route knowledge. A presentation of a survey map in the homepage of a document provides survey information at the outset as well, but in a different form than route knowledge.

A Question of Scale

If physical space is cognitively related to hyperspace, how many nodes does a hypertext document need to have to approximate the mental space equivalent to the size of a neighborhood or a city or province? In the real world, we have different cognitive strategies for dealing with different sizes of spatial configurations. How well hypertext cognitively maps onto physical space may be dependent on appropriate scale matches.

1.5 Hypertext Navigation

To this point, this chapter has focused on a broad overview of some of the issues regarding the spatial metaphor. The following section presents specific studies focused on the contribution of spatial navigational aids to hypertext performance.

How can we ascertain people's conceptualizations of such an intricate environment? As in most research on human cognition, such study requires inferring cognition through people's observable interactions with specific aspects of an environment. For hypertext research, a promising approach is to examine the influences of tools intended to facilitate the user's completion of a given task. Navigation tools can be used as mechanisms for closely examining the users' reliance on certain types of information to make navigation decisions. Navigation tools are easily manipulated to have the content and structure desired by the researcher. They can be designed to contain specific combinations of information. Navigation tools are often the users' only exposure to all parts of the document in a single format.

Information access in hypertext documents is impacted by the structure of the document itself outside of the navigational aids that may be provided (Korthauer & Koubek, 1994). There are three levels at which this may occur: 1) The inherent structure of the material (Is hierarchical organization an intrinsic quality of the material regardless of presentation?), 2) The organization of the material within the document or document structure cues (Are pages from subordinate categories differentiated from pages in superordinate categories?), 3) The linking structure (Does it reflect both of the above or, in and of itself, does it impose structure or constraints on the document?). These are global features of the hypertext document that are beyond the focus of this thesis. However, we need to be aware of and explicitly report these features during the course of investigating navigational aids.

1.5.1 Measures

How is navigational performance typically measured? Two categories of measurement encapsulate most of the measures typically used in hypertext research: efficiency and effectiveness measures.

Efficiency measures are based on speed and the number of steps taken to complete an information search. Efficiency measures often include the number of nodes visited in relation to the entire document, the number of excess nodes visited, the number of nodes revisited, the time taken on each node of the document and the total task time.

Effectiveness measures focus on the user's search accuracy as well as his/her recall and understanding of the structure of the document. Effectiveness measures typically include measures of recall of the elements within and the structure of the document. This would include the number of node titles remembered, number of node titles filled into a blank overview or a sketch map of the document (Stanton, Taylor & Tweedie, 1992). Accuracy is also an effectiveness measure denoted by the number of correct answers to questions that require the user to conduct an information search.

A distinction needs to be made between navigation effectiveness and learning of the document material. In contrast to effectiveness, learning involves a deeper understanding of the content of the document and is often measured by the study or review time per node, the number of main ideas or concepts retained (Dee-Lucas &

Larkin, 1997), and the number of concepts retained over time (McDonald and Stevenson, 1999).

Several studies that have found positive effects for spatially based navigational aids have reported an increase in efficiency measures but not in effectiveness measures. (McDonald & Stevenson, 1998, 1999; Wenger & Payne, 1994). If hypertext is to be a truly useful informational tool both measures would seem important. Making decisions quickly and accessing the least windows necessary is relevant but if users are not able to find the information they need (incomplete or inaccurate search) or recall how they found it, the usefulness of hypertext could be severely constrained. Therefore, it would seem that a successful navigational aid should increase both efficiency and effectiveness.

1.5.2 Types of Navigational Aids

A spatial map, spatial overview or graphical browser is a visual representation of the structure of the document, which depicts the hypertext nodes and the links that connect them (McDonald & Stevenson, 1998, 1999; Dee-Lucas & Larkin, 1995). These are usually in a diagrammatic form such as block diagrams, diagrams organized around a central term (spider map), or hierarchically ordered tree diagrams.

Non-spatial navigational aids or textual aids include alphabetical or arbitrarily ordered content lists, conceptually ordered content lists and indices (e.g., Simpson & McKnight, 1990). It is difficult to pinpoint what features distinguish a textual from a spatial aid. I suggest that lists that assign a spatial location through features such as indentation (e.g., Chou & Lin, 1997) should be thought of as a spatial aid as they provide a spatial location for each term in the layout.

1.5.3 The Effects of Navigational Aids

Most hypertext tasks can be described within the broad classification of open or closed tasks (Marchionini, 1989). Open tasks have a general goal of browsing or exploring the document. Closed tasks have a specific goal of finding a particular piece of information. The effects of navigational tools on both types of task have been studied. The following section describes first a general meta-analysis on navigational aids, task types and user characteristics, followed by specific examples of studies on closed and open tasks.

Chen & Rada (1996) conducted a meta-analysis on 20 experimental studies and 3 doctoral dissertations that appeared between 1988 and 1993. Generally, the meta-analysis indicated that the effects of navigational tools such as indices, table of contents and graphical maps interact strongly with the type of task involved (open or closed).

More specifically, for effectiveness scores (based on various measures of accuracy), task complexity showed the greatest effect size followed by graphical maps. For efficiency scores (based on various speed measures), task complexity and spatial ability showed the largest effects. Also, spatial ability interacted with navigational aid on measures of efficiency (speed).

These results suggest that spatial processing may be play a role in hypertext processing. However, as a meta-analysis is an analysis of common features across studies, it does not make detailed comparisons of the specific features of each type of navigation tool. It is still not clear which specific features may contribute to this effect.

Effects of Navigational Aids on Closed Tasks.

McDonald and Stevenson (1999) reported on two studies examining the effects of navigational aids on navigation and learning. In the first study, the effect of a localized spatial map, a textual content list and no aid were compared. The map consisted of labels with connecting lines indicating the links between nodes. Not all nodes were linked. The content list was a scrollable list of all the nodes in the document. The navigational aids were not interactive in that users could not access pages via the aids but had to navigate through the pages indicated by the map or list. The aids were displayed in separate windows from the text content of the document. Subjects were asked to first find ten target nodes and immediately after were asked twenty questions regarding information in the text. The navigation measures were: the mean time taken to access the target nodes and the number of additional nodes opened beyond that determined to be the optimal route. The learning measures were: the number of questions answered correctly and the mean number of node titles recalled during free recall. The results show that subjects in the map condition outperformed both the content list and no aid conditions on both measures of navigation efficiency. However, no significant differences were found between the three conditions on accuracy on the test questions. The number of titles recalled were comparable between the spatial map and the contents list but were significantly less for the no aid condition. The spatial map did not produce better recall than the content list.

In McDonald and Stevenson's second study, a spatial map and conceptual map were compared to a no aid condition. The conceptual map was similar to the spatial map described for the first study but with the added feature of link descriptions. The same measures were recorded. Subjects again found ten target nodes and but then answered forty questions based on the text, twenty factual questions and twenty questions meant to tap into deeper conceptual learning. This session was referred to as the acquisition phase. A week later, subjects returned to answer forty additional questions to test for long-term retention (the retention phase). As in the first study, significant differences occurred between all conditions on both efficiency measures. On the effectiveness measures, there was no difference between the spatial and conceptual map conditions on the accuracy of factual questions in the acquisition phase. Both were better than the no aid condition. In the retention phase, conceptual map subjects answered more factual questions correctly than the spatial map and no aid condition. Regarding the conceptual questions, in the acquisition phase, subjects in the conceptual map and no aid conditions answered more questions correctly than the spatial map subjects. In the retention phase, there was a difference between all three groups, conceptual map subjects answering correctly most often followed by the no aid group and finally the spatial map group. These results suggest that a conceptual map is important for learning whereas a spatial map increases efficiency but not effectiveness.

Dias and Sousa (1997) studied the effects on performance of a navigational map, described as a content index that provides a global overview of the document as well as a representation of the path taken by the user. Subjects were asked to find the answers to 12 questions. A map icon allowed the users to access the overview. From the overview, the subjects could access any page within the document. Performance was measured on the accuracy of the information search (answers to the test-questions), time spent on each screen page, frequency of visits to each information node, frequency of visits to the map and time spent on the navigational map. Performance on these measures was compared to a pre-defined, optimal path for that particular test item as determined by the authors. Subjects were divided in to three groups according to map access, both frequency and time spent on the map. The results showed no significant differences between the three groups. Subjects who visited the map frequently and /or spent more time on the map did not access more relevant screen pages or fewer unnecessary pages than subjects who used the map less frequently or for less time.

In another search task study, Leventhal, Teasley, Instone, Rohlman and Farhat (1993) found that measures of navigational performance were not correlated with use of a hierarchical overview. Correlations between accuracy and speed measures and both the percentage and total number of visits to two types of hierarchical overview cards were not significant.

Stanton, Taylor and Tweedie (1992) found evidence against facilitation of navigation maps on effectiveness measures. Subjects were divided into map and no map conditions and were asked to search out the answers for a sentence completion task. They were then asked to draw an outline of the hypertext document on paper. The results indicated that the no map subjects outperformed the map subjects on accuracy of the sentence completion task. Subjects in the no map condition were also better able to produce a drawing of the document structure and reported feeling more in control of the search situation.

Effects of Navigational Aids on Open Tasks.

The category of "open tasks" encompasses a number of different approaches. These have included: 1) directed reading of the document followed by free review, 2) reading until the users felt all nodes of the document had been accessed and 3) navigation of a document with no specific instructions. Such studies have produced mixed results.

Dee-Lucas and Larkin (1995) investigated the effects of structured versus unstructured overviews on recall. The structured overview (described to the participants as a content map) consisted of a hierarchical diagram with one central heading and three levels of subheadings. Through the placement of the headings, referred to as unit titles, the diagram indicated the subordinate and superordinate relations between the unit titles. The unstructured overview (described to the participants as a content list) was a vertically arranged list of each unit title. The unstructured overview provided no indication of the relations between the unit titles. Both overviews were interactive, in that, as the reader finished a unit he/she was required to return to the overview to select and link up to a new unit title. Both document formats contained 9 units. Participants were asked to read all material first in a specified order and then were free to review any material they chose. Participants were then tested on both recall and study strategy measures. The recall measures included the number of unit titles recalled, the proportion of text propositions (main ideas) recalled, the proportion of propositions recalled per unit and the proportion of units with at least one proposition recalled. On all these recall measures there were no significant differences between the two overviews.

Wenger and Payne (1994) investigated if the presence of a graphical browser impacted the comprehension and retention of the content and structure of a hypertext document. In this study, the graphical browser indicated the structure of the document in a hierarchical fashion and also, through the highlighting of visited nodes, allowed the user to see which portions of the document had been accessed. To minimize imposing a certain order of viewing, users were allowed to access the document at any entry point by choosing, on the first node, from an unstructured alphabetical listing of the contents. Both the browser group and the no-browser group had this option. The browser group could return to the browser from any node in the document. The subjects were instructed to read through all the nodes and indicate when they felt they had accessed all available nodes. At the end of the text, in a timed recall task, users were required to note as many as the node titles as they could remember. They then completed multiple-choice tests as an indicator of their comprehension of the text. Finally, to tap into the users' knowledge of the text structure, users were presented with pairs of node titles and asked to state if the nodes were linked within the previous text. On all three of these effectiveness measures, results showed no differences between the browser and no-browser groups. The only significant differences occurred in efficiency measures: the number of nodes accessed and the number of nodes revisited. Subjects in the browser groups accessed a higher proportion of the total available nodes and exhibited less revisitation than the no-browser subjects. Subsequently, the authors suggested that presentation of a graphical browser enhanced the efficiency of browsing but did not improve users' performance on measures of effectiveness.

Generally, these prior studies on the effects of navigation aids have provided inconclusive evidence on the role of conceptual vs. spatial information.

1.6 Summary of Review

Research concerning hypertext use is conducted within numerous disciplines each with its own focus, methods and terminology. Such results are reported in diverse sources making the integration of knowledge and theoretical comparison challenging.

Understanding hypertext use should not be confined to understanding this specific context or process but should also address connections to environments that share characteristics (e.g., real space) and to broader generalities of human cognition (e.g., spatial cognition, metaphor use).

The spatial metaphor suggests that spatial information and processing plays an important role in hypertext navigation but empirical evidence to date has not specified the parameters of that role.

1.7 The Current Research

Since these early discussions regarding the spatial metaphor, some would argue that the metaphor has subsided in its prevalence because the use of spatially-based tools has diminished. However, because spatiality is such a pervasive, unavoidable aspect of our environment and our cognitions, others remain convinced that understanding the role of spatiality in hyperspace is critical.

" It is our contention that an essential element in understanding ICTs (information and communication technologies) and cyberspace is a comprehension of how they are transforming, and creating new, spatialities , spatial forms and space-time relations. In short, we argue that geography continues to matter, despite recent rhetoric claiming the 'death of distance'. (Dodge & Kitchin, 2001)

Hence, a distinction should be made between the potential usefulness of the spatial metaphor itself as a psychological foundation for hypertext navigation versus the usefulness of the tools that have been developed thus far to support the metaphor. It is my assertion that a more thorough understanding of the cognitive processes implied by the metaphor may ultimately lead to a better applied use of the metaphor. The intent behind this thesis is to provide a theoretical investigation of the underlying cognitive processes related to the spatial metaphor rather than to present applied outcomes that directly address design considerations. This thesis investigates hypertext navigation from a cognitive psychology perspective with a focus on the form of mental representations that support this task. The navigation tools used within the human data phase of the thesis function as tools for examining cognitive processes rather than as applied solutions for navigation difficulties. If spatial processing is an integral part of hypertext navigation then providing spatially-based navigation tools should enhance the navigation process including the formation of appropriate mental representations of the document space. Given that spatiality can be represented in a multitude of ways within a navigation tool, this thesis should be viewed as a preliminary investigation to be followed by studies portraying spatiality in a variety of presentations.

The thesis begins with what could be referred to as a coarse-grained assessment of overt navigation behaviours followed by a series of increasingly fine-grained analyses aimed at examining underlying mental representations.

The first line of inquiry (Chapter 3) examines the impact of different types of navigation tool information on basic user behaviours such as the speed and economy of navigation during an information search task. Measures of accuracy and recall are also taken. The goal of this analysis is to ascertain if spatial information plays a unique role in influencing hypertext performance. If, as the spatial metaphor suggests, spatial processing is essential to navigation in hyperspace, navigation tools with higher degrees of spatial information should result in more efficient and effective navigation performance.

The second analysis of navigation behaviours (Chapter 4 – The Behavioural Space) examines the path patterns of users to determine if the presence of different types of information influence the sequence of users' navigation decisions. This chapter highlights the innovative application of a statistical technique (multidimensional scaling) that provides a rich but understandable format of navigational behaviour. The goal of this analysis is to better understand the relationship between properties of the navigation tool and properties of users' paths. The MDS configurations provide a means to make qualitative comparisons between tool and path characteristics and also provides an opportunity to uncover potential contributing dimensions that are not initially obvious.

The third analysis (Chapter 4 – The Representational Space) delves deeper into the issue of mental representation through the use of a distance-like ratings task. Users were asked to rate the relatedness or closeness of pairs of hypertext pages from the document they traversed. These ratings should reflect the users' underlying mental organization of the document space. The goal of this approach is to detect relationships between the navigation tool information and the users' mental representation of the document. This relationship may not be direct but rather could be mediated by the course of navigation behaviours from the beginning to the completion of the information search task. Therefore, relationships between behaviours and ratings are also examined using correlation and regression methods.

Finally, Chapter 5 provides an indepth exploration of spatial mental representations through the use of artificial neural networks. The goal of this study is to better understand the relationship between the form of mental representation and consequent behaviours. Two parallel distributed processing networks are trained to complete a spatially-based task. The internal representations developed by the networks to accomplish the task are carefully inspected by using correlations between network activations and actual distance information. This investigation provides insight into correspondences between overt behaviours and covert mental representations. The networks reveal that the properties of underlying representations are not always equivalent to the behaviours they produce. This finding underscores the necessity for understanding what type of mental representations are created by given presentations of information such as those produced by particular navigation tools. Relevant to the topic of hypertext navigation, without closer inspection, it cannot be assumed that spatial navigation tools create spatially-based mental representations.

CHAPTER 2: TRACKING COGNITIVE PROCESSES IN HYPERTEXT

2.1 The Data Collection Program

The task of navigating in hypertext involves the integration of many complex cognitive processes within a very multifaceted environment. The multitudes of variables that such a task presents demands methodologies that allow for control of the interface that users see and that provide flexibility for collecting diverse types of data. To provide such control, the human data used in this thesis was collected using a customized program. The purpose of the following section is to describe in detail the computer program developed to record behaviours associated with hypertext navigation.

2.1.1 Controlling the Interface

The Visual Basic ActiveX WebBrowser Control included in Visual Basic 6.0 creates a highly customizable web browser that allows for the management of options available to users of hypertext documents. This is important for two reasons. First, to access information about the cognitions involved in hypertext use, a pared-down, sparse version of browsers and hypermedia pages may be desirable. Custom programming allows the researcher to introduce new elements to the browser controls as well as eliminate unwanted controls resulting in interface elements that are specifically applicable to the question of interest. For example, for some research questions, the elimination of options such as the back, forward or print buttons from the browser menu may be favorable.

Second, a custom browser ensures a novel environment for every user regardless of their prior experiences with standard browsers. Pretesting users on previous computer use is helpful but these measures usually rely on general information such as how often a computer is used in a certain timeframe or the number of activities the user engages in on a regular basis (e.g., e-mail, word processing). Questionnaires can result in ambiguous responses depending on the questions asked and the users' interpretations of the questions. Ensuring that every user begins the research task with an equal level of familiarity with the browser eliminates one extraneous variable.

2.1.2 Collecting Appropriate Data

Often, hypertext researchers are interested in tracking very quantitative navigation responses such as time spent on pages, number of pages accessed and order of pages. A custom program can provide procedures for data collection that are strictly controlled by the researcher, eliminating any impact of features inherent in other standard tracking mechanisms such browser history lists or server logs.

For example, browser history lists make use of a variety of procedures for recording page access. To optimize the history list, some methods bump prior revisits and only retain the most recent instance of access (e.g., for further description of history lists see Tauscher & Greenberg, 1997). By creating a custom application, a local log file of users' movements can be constructed to record every single move with no elimination of repeated page accesses that can happen with the typical history lists of standard browsers.

For some types of data, collecting data via server logs presents obstacles as well. The procedures used by the server to cache items may skew data by affecting the way pages are delivered to the user. For instance, as a single page is accessed numerous times, by numerous users, the server retains the page in memory and will load that page quickly relative to a page that is not held in memory. This results in navigation experiences that may be different from user to user. Also, because server use involves multiple users on a shared network, system traffic could interfere with timed results as page delivery could be delayed by network congestion. By having tracking procedures local to each individual workstation, that is, creating local logs instead of server logs, we can avoid the above-mentioned problems.

2.1.3 Advantages of Visual Basic 6.0

It is true that other programming formats include a customizable browser function, so why use Visual Basic 6.0? There are several pragmatic reasons for using Visual Basic specifically. First, it's easily accessible. Versions of Visual Basic are included in popular office suite software packages. Second, Visual Basic is well–supported in that many reference books are available as well as extensive on-line help. Third, Visual Basic is a relatively simple programming environment, as compared to many other programming languages.

The program description is organized into two sections. First, the development of the interface of the document is described. Second, the coded routines that accompany the interface are detailed.

2.1.4 Interface Development

The hypertext interface is created in two parts: 1) the actual hypertext pages with the specified content and links are constructed and 2) the Visual Basic application forms, including the WebBrowser Control that displays the Web pages, are programmed.

Chapter 2

Developing the Visual Basic Application

Visual Basic 6.0 is an Integrated Development Environment (IDE) that allows the user to design, run, test and debug Windows applications built using the Visual Basic language. Development in the Visual Basic 6.0 environment involves placing standard components onto a visual container called a form, thereby constructing the user interface. Coded routines are created to respond to the user's interactions with the form's interface components. A user interaction is called an event and the coded response is called an event handler. Each control has a certain set of events that will generate responses (e.g. Mouse Click, Resized, Document loaded). Several forms, which appear as separate windows, were constructed for this application, including the form that contains the specialized component, the WebBrowser control.

The Development of Forms

Individual forms for the interface are constructed to accomplish certain tasks. Tasks are subdivided into user actions that are represented as various controls in strategic positions on the form. Each form in the application is maintained in a file designated by the form name plus the extension.frm. Form files contain the graphic interface of the form but are also text files that contain descriptions of the controls on the form and coded event handlers.

There are several types of standard window controls that can be included on a Visual Basic form. For example, Textboxes are used to display and enter text. Command buttons represent actions that will occur when that control is activated by the user. Option buttons are presented in groups and are used when the user is required to make one choice out of several options. The CheckBox control allows the user to make several choices from a list of options. In Figure# 1, the Login form shows an example of a form with labeled textboxes (the spaces provided for user and password submissions), option buttons (the buttons used to select the appropriate condition) and a command button (the "submit" button). Such controls are added by selecting and placing icons from the General Toolbox panel onto the form. Numerous properties can be selected for each control (e.g., for the textbox control, properties for the number of text lines, the inclusion of scroll bars and the maximum length of the entered text can be chosen). Properties can either be set from the Properties window or written into the code. As the project is developed, the Project Explorer provides a list of the components (e.g., forms) that have been added to the project.

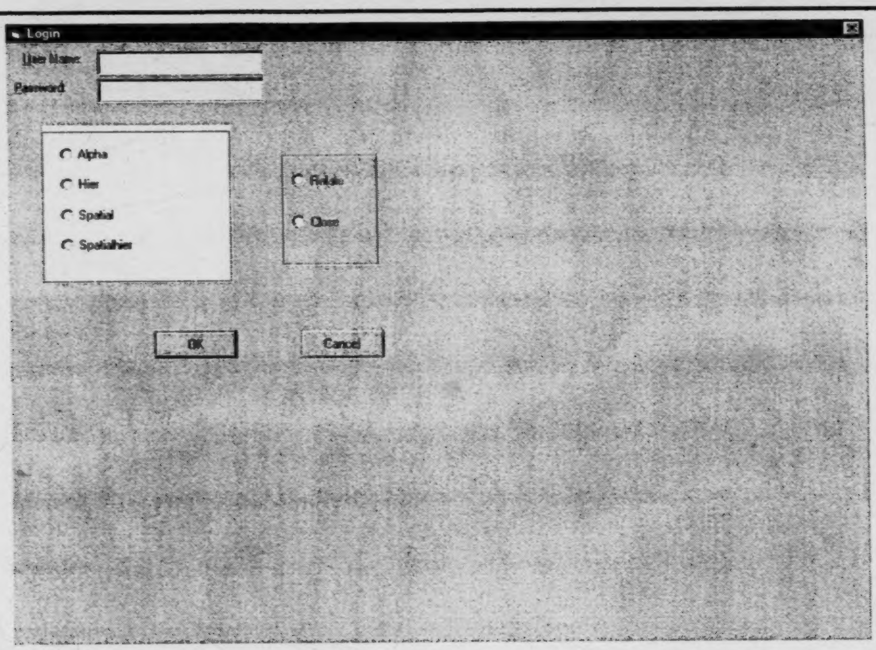

Figure 2.1. Administrative Login Form

2.1.5 Code that Accompanies the Interface

Controls can be activated by several types of events triggered by user interactions in the form of either mouse use (e.g., clicking, press and release of the mouse button or movement of the mouse) or keyboard use (e.g., key press, changes to text boxes). Other events are triggered by changes in control state (e.g., the WebBrowser's documentComplete is triggered when the control has successfully loaded a page). Whenever a control on a form is activated, the program will initiate the event handler routines that are coded for the event that activated the control. For example, on the Login form (Figure # 2.1), once selections are made between the option buttons in the dialog boxes, clicking on the "submit" button initiates the subroutine coded for the on click event as outlined in Listing # 1(see Appendix A). Note the initial section of code for each form declares the variables local to that form using the *Dim* (dimension) statement. Local variables that are intended to be available only to a specific subroutine can also be declared within individual subroutines.

An initial login form, such as the one above, can be used to specify a particular condition or version of stimuli for each individual subject. Later in the application, flow control statements (e.g., "If-Then", "Select Case") can access the selections made in the initial Login form and use them to execute only the appropriate statements for the specified stimuli.

Displaying the Web Pages

In this application, the Browser Form contains all the code necessary to display the Web pages. When the Browser Form is loaded, the "select case" structure is used to direct the program to present the appropriate instruction page and the navigation aid to be used in that specific test session (Listing # 2, Subroutine #1 and #2 from the Browser Form set-up routines, see Appendix B and C). Within the code, the four conditions in this study were referred to as Alpha, Hier, Spatial, Spatialhier.

Once the WebBrowser Control is functioning, the links embedded in the individual Web pages allow the user to revisit the navigation aid and instructions page. Each successful navigation triggers the WebBrowser's "Document Complete" event, which causes the subroutine WebBrowser1_DocumentComplete to execute, calculating and recording the time the user has spent on the previous page (Listing #3, Subroutine #5, #6, see Appendix D).

Developing Data Files

In order for the program to open data files and store data within them, certain variables must be accessible to all Forms and Modules throughout the run time of the application. These variables must be declared as *Public* (rather than *Dim*) and as Form variables or within Modules.

Users' responses to the four tasks were collected and stored in text files that are automatically constructed when the user interacted with specific forms in the program. For three of the tasks this merely involved transferring the typed responses or option button and checkbox responses of the user to text files. The fourth task required the storage of selections of text, which necessitated a slightly different approach. In order to check the accuracy of the users' search efforts, users were instructed to highlight the answer with the mouse (e.g., phrase or sentence) for each question from the page on which they found it and then submit the answer by clicking on a "submit" command button. The code that responds to the "Click" event takes advantage of the fact that the WebBrowser is an OLE container for the active HTML page and as such can access the active document within the container. The container accesses the document through an interface that initiates the correct implementation of the common command specified. One of the commonly supported commands is to copy the currently highlighted or selected text to the clipboard. Using the execCommand method of the active document's interface with "Copy" as the indicated command, the highlighted or selected data is copied to the clipboard. The globally available (common to and useable by all) clipboard was then accessed by the research program, which retrieved the data to a local variable and wrote the information to file (Listing # 3, Subroutine #3, #4, see Appendix E). Constraining the user to copy and pasting information to the clipboard avoids common problems like typing mistakes or user interpretation of data, which would make subsequent data verification or validation difficult.

The collection of efficiency measures was accomplished using a log file of the users' movements through the document which records each page accessed including revisitations, the time spent on each page and the page where the user found the information required by the search task (see Figure # 2.2 an example of the log file). The log file was constructed by writing successive lines in the log file each representing a record of specific structure (e.g. User, Page, Time Spent). Each record is a series of comma-delimited pieces of information. The new line character separates individual records. The format of the log files was chosen because it offers a format that is tool independent providing guaranteed access while being of a common format used by data analysis tools that import data from other sources. This allows for easy organization of data beyond the data collection phase.

```
"Subject 1   ","C:\VB Project2\Instructions.htm","0.719999999993888"
"Subject 1   ","C:\VB Project2\instructions1.htm","47.5800000000017"
"Subject 1   ","C:\VB Project2\SpatialHier.htm","19.4700000000012"
"Subject 1   ","C:\VB Project2\Shtypes.htm","25.9100000000035"
"Subject 1   ","C:\VB Project2\SpatialHier.htm","4.63999999999942"
"Subject 1   ","C:\VB Project2\Shclub.htm","18.6100000000006"
"Subject 1   ","C:\VB Project2\SpatialHier.htm","1.44999999999709"
"Subject 1   ","C:\VB Project2\Shsac.htm","23.9499999999971"
"Subject 1   ","C:\VB Project2\SpatialHier.htm","1.68000000000029"
"Subject 1   ","C:\VB Project2\Shconjugat.htm","23.0600000000049"
"Subject 1   ","C:\VB Project2\SpatialHier.htm","1.83999999999651"
"Subject 1   ","C:\VB Project2\ShChytrids.htm","26.1800000000003"
"Subject 1   ","C:\VB Project2\SpatialHier.htm","1.13999999999942"
"Subject 1   ","C:\VB Project2\Shlichen.htm","24.5999999999985"
```

Figure 2.2. Example of a Log Text File.

The Visual Basic program outlined above provides an accessible and effective means of tracking the processes involved in hypertext use. Once the custom browser program is created, altering the content and functions of the individual hypertext pages can modify the task to address additional research questions. To accompany navigation data, relevant tasks such as recall and recognition tasks can be programmed for on-line data collection in a text format that is easily manipulated for subsequent data analysis.

2.2 Methods

The portion of this thesis using human data focused on the reliance of users on different types of information within a navigation tool. Four different versions of a navigation tool were presented to four groups of participants. Participants were required to complete a series of tasks involving the same hypertext document, the only difference being the navigation tool presented. The following section provides details for the methods used for this part of the thesis.

2.2.1 Participants

The participants were 169 undergraduate first-year psychology students (81 males and 88 females) from the University of Alberta, Edmonton, Alberta, Canada. Students participated as part of a research participation program and received credit for their participation. The students were randomly assigned to four conditions based on the navigation aid presented: Alphabetical, n = 44 (22 males and 22 females), Hierarchical n = 41 (20 males and 21 females), Spatial, n = 41 (18 males and 23 females), and Spatial/Hierarchical, n = 43 (21 males and 22 females).). Participants were pretested on their Internet experience via a questionnaire. Items on the questionnaire included the frequency and duration of Internet use in a one week period, the availability of Internet access at home, and a rating of how much each participant engaged in particular Internet activities such as information searches, surfing, chat rooms, e-mail, use of search engines, copying images and the construction of a Web page. Preliminary analyses found no differences between the four groups on Internet experience.

2.2.2 Materials

The Hypertext Document

Participants were tested on a 22-page hypertext document on the topic of Fungi programmed using Visual Basic 6.0 (Professional Edition) (Boechler, Dawson & Boechler, 2002). Each hypertext page contained a short text section with similar word counts between pages (approx. 130 words each) and a picture to help users differentiate between pages (see Figure 2.3).

The Navigation Tools

The four navigation aids used in this study were constructed to approximate the most singular representation of each type of information, the absence of both types of information or a combination of both types of information. The Alphabetical navigation tool is merely a content list of the page titles organized alphabetically. It contains minimal spatial information and no conceptual information, as the list provides no suggestion as to how the page titles may be semantically related (see Figure 2.4a).

The Hierarchical tool contains conceptual but minimal spatial information. It is a content list that designates the superordinate and subordinate categories in the material through the use of font size and label color (Figure 2.4b).

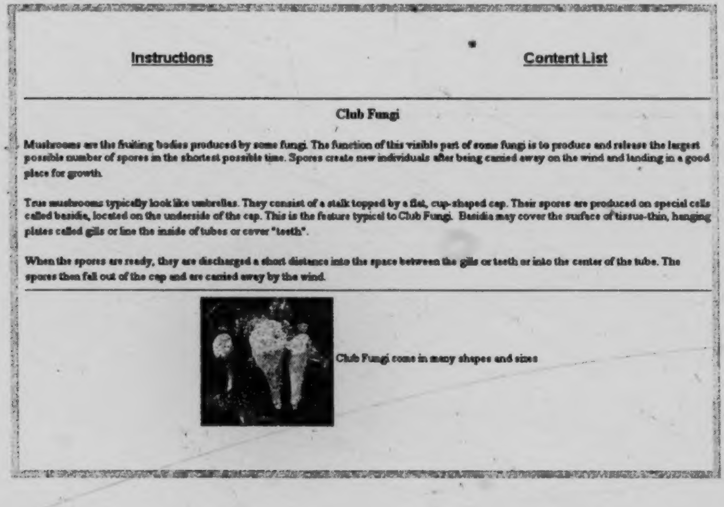

Club Fungi

Mushrooms are the fruiting bodies produced by some fungi. The function of this visible part of some fungi is to produce and release the largest possible number of spores in the shortest possible time. Spores create new individuals after being carried away on the wind and landing in a good place for growth.

True mushrooms typically look like umbrellas. They consist of a stalk topped by a flat, cup-shaped cap. Their spores are produced on special cells called basidia, located on the underside of the cap. This is the feature typical to Club Fungi. Basidia may cover the surface of tissue-thin, hanging plates called gills or line the inside of tubes or cover "teeth".

When the spores are ready, they are discharged a short distance into the space between the gills or teeth or into the center of the tube. The spores then fall out of the cap and are carried away by the wind.

Club Fungi come in many shapes and sizes

Figure 2.3. Example of hypertext document page

The Spatial tool was a map of the page titles containing spatial information but minimal conceptual information (Figure 2.5). Given that spatial information cannot be completely eliminated from any type of visual display, even the content lists contain some spatial information as labels can be allocated a location (e.g., second from the top), there exists an infinite number of possible spatial arrangements that could be used for the spatial version of the navigation tool. As a preliminary investigation into the role of spatiality in navigation tools, I operationalized the formatting of the spatial navigation tool in the following manner: The spatial array was derived by counting the number of steps between titles in a hierarchically ordered tree diagram of the material. The resulting "distances" were transformed into a matrix of correlations that was analyzed using multidimensional scaling. The configuration produced by this analysis provided the positions for the page labels on the map.

Anti-rejection Drug
Biocontrol
Caterpillar Tonic
Chytrids
Club Fungi
Conjugation Fungi
Disease
Food
Fungi
Heart
Lichen
Medicine
Morels
Negative Aspects
Penicillin
Poisonous Mushrooms
Positive Uses
Sac Fungi
Skin
Truffles
Types
Yeast

FUNGI
POSITIVE USES
Medicine
Penicillin
Anti-Rejection Drug
Caterpillar Tonic
Food
Morels
Truffles
Yeast
Biocontrol
TYPES
Club Fungi
Sac Fungi
Conjugation Fungi
Chytrids
Lichen
NEGATIVE ASPECTS
Diseases
Skin
Heart
Poisonous Mushrooms

Figure 2.4. a) Alphabetical List b) Hierarchical list

MDS is a statistical method for uncovering the structural regularities or patterns hidden in a matrix of data and representing that structure in graphical form for easier visual interpretation. The coordinates that create the MDS configuration represent some underlying property of the data and the configuration of points they create illustrates how closely the objects under study are related to each other as a function of this underlying property. In the case of the spatial tool, the coordinates reflect the distances between labels on the tree-diagram, a diagram that has explicit hierarchical categories.

In an MDS analysis, the dimensions that produce the spatial coordinates may be abstract such as the conceptual relatedness between pages or concrete such as the physical distances between page labels on the navigation tool. On a MDS plot, related items are plotted close together and unrelated items are plotted far apart. (for examples see Borg & Groenen, 1997). In the case of this MDS configuration, both conceptual relatedness and physical distance are contributing dimensions as these are salient factors in the tree-diagram arrangement.

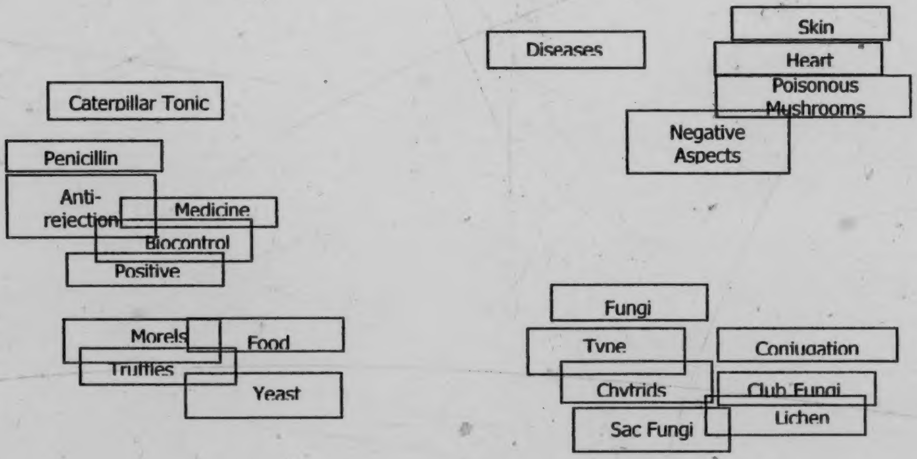
Diseases
Skin
Heart
Poisonous Mushrooms
Negative Aspects
Caterpillar Tonic
Penicillin
Anti-rejection
Medicine
Biocontrol
Positive
Fungi
Type
Conjugation
Chytrids
Club Fungi
Lichen
Sac Fungi
Morels
Food
Truffles
Yeast

Figure 2.5. Spatial Navigation Tool

The spatial map did not provide explicit information about the semantic relatedness of pages but the process described above created a configuration with some implicit information about relatedness given that the clustering of page titles reflected to some degree the conceptual levels of the tree diagram. This approach was chosen over totally random spatial label locations as it is unlikely that a hypertext document would be accompanied by a random spatial map as a navigation aid. To avoid any explicit indications of semantic relatedness, link designations were excluded from the spatial tool.

Throughout the thesis, this spatial array will be referred to as the spatial navigation tool. This does not preclude the possibility that other spatial arrangements may produce different results. Future research needs to investigate the effects of many different versions of spatiality.

The Spatial/Hierarchical tool explicitly contains both types of information. It is a tree-diagram with the superordinate heading of "Fungi", below which are three levels of subordinate categories (see Figure 2.6). This tool presents conceptual information not only through the levels of categories represented but also through the link connections between subsets of labels. Link connections only occur within the three depth levels of the hierarchy. The tool also contains spatial information as the location of labels corresponds to their conceptual relationship to other labels. Labels that are closely conceptually related are closer together in the configuration whereas labels with less semantic relatedness are farther apart. As with the spatial tool, a variety of spatial configurations could be used. This particular configuration corresponds to the inherent hierarchical organization of the topic of "Fungi".

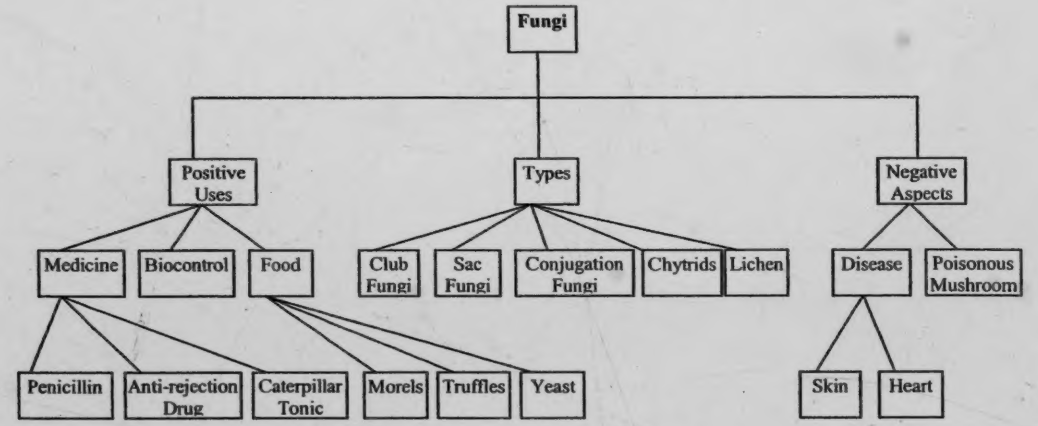
Fungi
Positive Uses
Types
Negative Aspects
Medicine
Biocontrol
Food
Club Fungi
Sac Fungi
Conjugation Fungi
Chytrids
Lichen
Disease
Poisonous Mushroom
Penicillin
Anti-rejection Drug
Caterpillar Tonic
Morels
Truffles
Yeast
Skin
Heart

Figure 2.6. Spatial/Hierarchical Navigation Tool

The organization of the navigation aids did not reflect the linking structure of the document but only served as a representation to help the users mentally organize the document material. All page labels on all four navigation aids were linked to the corresponding page, therefore, the users could always navigate directly to any page from the navigation aid without having to access intermediary pages. Hence, the optimal route to a given page was always a direct selection of a particular page label from the navigation aid. These tools were designed with the recognition that there is no representation that is purely one type of information and also with the constraint of retaining some ecological validity.

The Tasks

Participants were asked to complete four tasks: an information search task, a free recall task, a recognition task and a distance-like ratings task.

Information Search Task: The search task began with an instruction page, followed by the navigation aid for a given condition. It consisted of ten questions presented consecutively on various aspects of fungi. Users were required to navigate the document and locate and record the phrase or sentence that answered each question.

Recall Task: The free recall task was intended to determine if there were differences between groups in how well participants could remember elements of the document content, specifically page titles. With the free recall task, participants are not cued for responses but must generate them independently. The recall task form appeared automatically after the response was submitted to the final question of the information search task. The form consisted of an instruction caption and a textbox for participants to type in the titles of pages they remembered. The instructions read, " In the text box below, type all the page titles you can remember from the "Fungi" document. When you are finished, click the "Submit" button at the bottom of the page".

Recognition task: The recognition task was again a measure of memory for document elements. As the page titles were single labels (one or two words), the recognition task was modeled after a title recognition task developed by Stanovich and West (1989) to test print exposure. It consists of a list of titles with both actual and foil titles. Participants are asked to indicate the titles they judge as familiar. In Stanovich and West's version, participants are told that some titles are actual titles and some are not and that they should, therefore, refrain from guessing. Participants are not told the ratio of actual titles to foils. In the current recognition task, the participants have already been exposed numerous times to the entire list of titles via the navigation tool. Therefore, it was assumed that the participants would recognize the list was substantially longer than the number of titles they were exposed to in the navigation tool and were not informed that foils were present. The current recognition task consisted of 22 page titles and 10 foils.

The recognition task immediately followed the recall task. The form consisted of an instruction text box at the top of the page and two columns of checkboxes with the page titles adjacent to them. The instructions for this task read, "From the list below, check all the page titles you can remember from the "Fungi" document. When you are finished, click the "Submit" button at the bottom of the page".

Ratings Task: The ratings task was designed to reveal how participants mentally organized the structure of the document. Immediately following the search task, participants were asked to rate 100 pairs of page titles according to their distance from or relatedness to each other. The form for the ratings task consisted of a text box with instructions at the top of the page followed by a list of page pairs on the left side of the page. On the right side, a series of buttons appeared with ratings of one to seven labeled above them. Each page contained fifteen pairs of page titles with the final page containing the remaining ten pairs. Two subgroups in each group were tested with different instructions. One group was instructed as follows "Based on a scale of 1 (lowest) to 7 (highest) rate the following A and B pairs according to the question: How related are pages A and B?”. The other group was given the instructions “Based on a scale of 1 (lowest) to 7 (highest) rate the following A and B pairs according to the question: How close are pages A and B?”. Participants used the mouse to select the button that represented their chosen response for each item.

2.2.3 Procedures

At the beginning of the test session, as a group, the participants were given brief instructions on the function of icons and the use of the mouse. Participants were then told they were required to complete four tasks related to a document that they would be viewing. It was explained that Task 1 would involve searching through the document for the answers to 10 questions and that once this task was complete the program would provide instructions for the following tasks. Participants were told that when the program proceeded to a Windows file folder window all the tasks were completed. Each participant was then assigned to a separate testing room.

To begin the information search task, participants were presented with a login page where they submitted their name and student number by typing in the necessary information into textboxes and clicking on a command button labeled “Submit”. This initiated the appearance of the instruction page, which students could view as long as they chose before clicking on a command button labeled "Continue". The navigation tool for that condition then appeared in a new window with the first question of the information search task listed at the bottom of the browser frame. The timer for the first question was initiated when the navigation tool page appeared. Also, on the bottom of the browser frame was a command button labeled “Found it”. When the participant had located the answer to the question they were to click on the "Found it" button and the timer would stop. The "Found it" label then changes to "Submit". When the student was sure of the answer he/she had selected and had highlighted the answer with the mouse, he/she would then click on the "Submit" button and the program recorded the answer. The " Submit" button was added to ensure that the recorded search time was not confounded by the time that participants might take to execute the selection of the phrase with the mouse and also to ensure that participants did not make impulsive or inadvertent selections that they could not alter before submitting them. The clicking of the "Submit" button also prompted the return to the navigation tool page and the appearance of the next question. This process was repeated until all ten questions were answered.

After the completion of the information search task, the program proceeded to the recall task. Students were instructed to type all the page titles they could remember from the document in the textbox below and then to click on the "Submit" button to record their answers.

Following the recall task, the recognition form appeared. Again, participants were to make their selections and then click on the "Submit" button.

Finally, the program continued to the ratings task. Participants were asked to complete the 100-item ratings task of page title pairs. On the completion of this task the program ended and returned to a Windows file folder window.

CHAPTER 3 - PERFORMANCE IN HYPERSPACE

3.1 Issues related to Navigation tools

This thesis focuses on examining the cognitive processes involved in hypertext use by manipulating the organization and content of navigation tools. This is a different goal from attempting to design the optimal navigation aid, as the navigation tools used in this study have limits that may make them impractical in an applied setting. Instead, I viewed these navigation aids as mechanisms for revealing the effects of certain types of information presented to the user. Specifically, I was interested in the effects of spatial and conceptual information in navigation aids on several measures: the efficiency and effectiveness measures described in the introduction (I refer to these as performance measures), the path patterns that users choose and the form of mental representation that users develop of the document. This current chapter focuses on the effects of different navigation tools on performance measures, specifically, speed, number of pages accessed, recall and recognition. The following data analysis addresses two issues introduced in previous research. The first a specific methodological issue and the second is a wider theoretical issue.

First, results of earlier studies on the influence of navigation aids on performance measures may have been clouded by the users' preference for accessing a particular type of tool over another type of tool. For example, McDonald and Stevenson (1999) found that when users were given the option of accessing a spatial map or conceptual map, spatial map users took less time to complete the task than conceptual map users. However, it was not reported how frequently each group accessed each navigation aid relative to other pages in the document. Dias and Sousa (1997) found that when users were allowed to access the map provided at will, different subsets of users accessed the map significantly more than other subsets. Other studies did not make clear if the users' only avenue for navigation was through the navigation tool (Dee-Lucas & Larkin, 1995, Beasley & Waugh, 1995). The current study forced users to make all navigation decisions from the navigation tool only, thus the impact of each type of information on performance and memory should be more salient as all users experienced similar exposure to the navigation aid. The total number of pages that an individual accessed might vary from user to user but the ratio of pages visited to navigation tool visits remains constant.

Second, one theoretical issue in hypertext research involves the role of spatial versus conceptual information in hypertext use. The design of the present study was intended to further separate the effects of these two types of information on navigation performance. An important aspect to note about a number of the previous studies comparing spatial navigation tools versus content lists is that hierarchical organization (conceptual information) and spatial location (spatial information) are sometimes confounded (e.g., McDonald & Stevenson, 1998, Dee-Lucas & Larkin, 1995, Beasley & Waugh, 1995, Wenger & Payne, 1994). For example, by placing page titles within a hierarchical framework, such as a tree diagram, pages are automatically assigned a spatial location producing a navigation aid that contains both types of information. Spatial maps, which focus on the spatial relation of pages, can also provide information about the conceptual relatedness of pages in several ways. First, individual pages can be linked or not linked. Second, the full linking of a subset of pages can create a cluster of related material versus the partial linking of other subsets. Third, the spatial clustering of page labels implies relatedness. Conversely, providing a textual content list removes not only spatial location but also the hierarchical information that highlights the conceptual relations between page titles.

Thus, some previous studies provide an incomplete picture of the contribution of spatial versus conceptual information to navigation performance. For example, McDonald and Stevenson (1998) found that a spatial map produced faster and more accurate searches in which users accessed less extraneous pages. However, this was in comparison to a content list that conveyed no conceptual information. Dee-Lucas and Larkin (1995) found that access to a spatial map resulted in better recall for the unit titles in a document than an alphabetical content list, which, again, does not relay conceptual information to the user.

In the current work, in order to understand more clearly the effects of spatial and conceptual information in hypertext navigation comparisons were made between four carefully constructed navigation aids. These aids and the procedures used to collect data with them were described in detail in Chapter two. Generally, within these four tools, spatial information was considered present when the page titles were organized in a configuration that was not a linear list, a list being the only representation that minimized spatial location as an element. Thus, page labels in the spatial tool were situated as items on a map, spread across the page surface. The spatial/hierarchical tool also situated labels throughout the 2D space of the page. Conceptual information was represented by the hierarchical organization included in two of the navigation tools, the Hierarchical and Spatial/hierarchical aids, which were organized around one superordinate heading and three levels of subordinate headings. This hierarchical organization conveyed information about the importance of particular page titles and the conceptual relatedness of the pages to one another.

3.2 Goals of the Performance Analysis

The goal of this set of analyses was to determine if the presence of spatial and/or hierarchical information in a navigation tool differentially influences the speed and economy of navigation as well as users' abilities to recall document elements and locate target information.

3.3 Results

On all the performance measures, preliminary analyses indicated no differences between genders. Therefore, tests for gender differences were not conducted within subsequent analyses.

3.2.1 Efficiency Measures

Effects on Speed of Navigation

Total task time: On the time measure, four subjects were lost due to missing data. Based on the criteria of greater than 3 standard deviations away from the mean on standardized scores, three data points were excluded as outliers. The resulting sample size for this task was n = 162. Significance for all measures was determined using a probability value of .05.

The means and standards deviations for total task time for each group are shown in Table 1. To detect any differences between groups in the speed of navigation, a one-way ANOVA (Time x Condition) was conducted on total task time (n = 162). There was a significant main effect for Condition, F (3,158) = 6.871, p < .001. Tukey's HSD post hoc tests were used to reveal where these differences occurred. (Alphabetical vs. Hierarchical, p < .01, Alphabetical vs. Spatial, p < .01 and Alphabetical vs. Spatial/hierarchical, p < .001).

The Alphabetical group performed poorly relative to the other three groups. The Spatial, Hierarchical and Spatial/Hierarchical groups perform comparably indicating that combining spatial and conceptual information did not speed up search any more than either alone.

Table 3.1
Means and Standard Deviations for the Average time across questions to complete the information search task for all groups.

	M	SD	n
Alphabetical	95.42	33.48	44
Hierarchical	76.09	25.60	40
Spatial	73.92	27.25	37
Spatial/hierarchical	72.06	19.07	41

Effects on Economy of Search

Number of pages accessed: On the page count measure, four subjects were lost due to missing data. Based on the criteria of greater than 3 standard deviations away from the mean on standardized scores, one data point was excluded as an outlier (n =164).

A one-way ANOVA was conducted on the mean number of pages accessed by each of the four groups (PageCount x Condition). Again, a main effect for Condition was evident, F (3,160) = 15.084, p < .001.Tukey's post hoc tests indicate the differences lie between the Alphabetical group and the other three groups (Alphabetical vs. Hierarchical, <.001, Alphabetical vs. Spatial, p < .001, Alphabetical vs. Spatial/Hierarchical, p < .001). Providing organization, whether it is spatial or conceptual, contributes to a more parsimonious search.

Table 3.2
Means and Standard Deviations for the Average number of Pages accessed across questions to complete the information search task for all groups.

	M	SD	n
Alphabetical	6.87	2.48	44
Hierarchical	4.39	1.53	41
Spatial	4.91	1.7	37
Spatial/hierarchical	4.67	1.68	42

3.2.2 Effectiveness Measures

Effects on Free Recall

Recall scores were obtained for all 169 participants. There were no outliers on the recall scores.

A one-way ANOVA was conducted on the mean number of page titles recalled by each of the four groups (Recall x Condition). Again, a main effect for Condition was evident, F (3,165) = 2.689, p <.05. Tukey's post hoc tests indicate this result was due to a difference between the Hierarchical group and the Spatial/Hierarchical group, p <.05.

The group with the added spatial information was more able to recall page titles than the group given only hierarchical information.

Table 3.3
Means and Standard Deviations for the Average number of Page Titles recalled for all groups.

	M	SD	n
Alphabetical	9.05	3.69	44
Hierarchical	8.66	4.41	41
Spatial	9.37	4.32	41
Spatial/hierarchical	10.89	3.13	43

Effects on Recognition

The recognition task produced complete data for 167 participants. The recognition score was calculated as the total number of correct titles chosen plus the number of foil titles not chosen. A one-way ANOVA (Recognition x Condition) revealed a significant difference between groups, F (3,163) = 20.177, p < .001.Tukey's HSD posthoc tests revealed this difference was due to the alphabetical group's elevated performance over the other three groups (Alphabetical vs. Hierarchical, p < .001, Alphabetical vs. Spatial, p < .05, Alphabetical vs. Spatial/Hierarchical, p < .001). Given that the Alphabetical group performed poorly relative to the other groups on several measures, this was a surprising result. Consequently, the format of the task was reviewed. It was found that the listing of the titles used in the recognition task was partially alphabetically ordered thus giving an advantage to the alphabetical group. Subsequently, a smaller sample of 48 participants was tested on a randomized version of the title list, which resulted in no differences between the four groups on the recognition scores.

Effects on Search Accuracy

Proportion of correct answers: Accuracy data was available for 166 subjects. Again, based on the criteria of greater than 3 standard deviations away from the mean on standardized scores, two data points were excluded as outliers (n = 164).

To detect any differences in the number of correct answers obtained by subjects in each group, a one-way ANOVA (Answer x Condition) was conducted on proportion correct. This analysis revealed a main effect for Condition, *F* (3,160) = 2.664, *p* = .05. Based on Tukey's HSD, this effect seems to be due to the differences between the Alphabetical group and the Spatial group (Alphabetical vs.Spatial , *p* = .063). However, the average scores for all four groups were high (see Table 4) indicating that the task was not difficult for most subjects. The questions used in this task were fact-based not concept-based. More challenging types of questions may have produced different levels of accuracy between groups. Evidence for differential effects on accuracy due to question type has been documented (Leventhal et al., 1993).

Table 3.4

Means and Standard Deviations for the proportion of correct answers from the information search task for all groups.

	M	SD	n
Alphabetical	.76	.26	43
Hierarchical	.84	.15	40
Spatial	.86	.14	38
Spatial/hierarchical	.85	.14	43

3.3 Discussion

One purpose of this analysis was to address the general issue of multidimensionality of navigation tools. In this instance, I was interested in the effects of spatial versus conceptual information on navigational measures, a question raised frequently in prior literature. However, directly measuring the differential effects of these two types of information is complicated by the fact that most navigation aids are not absolute singular representations of either type of information. Spatial maps often contain implicit conceptual information via the indication of links and the positioning of page titles relative to each other. Conceptual or hierarchically ordered diagrams assign a spatial location to labels even if their positioning is not related to the document material structure. The absolute degree of spatial versus conceptual information within or between navigation aids is not quantifiable. When comparisons are made between only two tools that are intended to reflect a singular representation of a particular dimension, the potential for delineating the effects of that dimension (e.g., spatial versus conceptual information) is diminished because the effects may be obscured or inflated by the presence of some other unspecified dimension. Therefore, in order to fully understand the role of spatial and conceptual information in hypertext use, it is necessary to separate the influences of each type of information. The four navigation aids in this study were carefully constructed to represent either a preponderance of spatial or conceptual information, a combination of both or the lack of both.

3.3.1 Summary and Implications.

Even in a relatively small hypertext document, we see differences in performance based on the information in the navigation tools provided. After constraining all groups to the same ratio of page to navigation tool visits, measures of efficiency and effectiveness showed few differential effects for spatial versus conceptual information. Generally, providing alphabetical information alone resulted in poorer user performance.

Viewing the two efficiency measures separately, for speed of navigation, as measured by time on task, either spatial or hierarchical information or a combination of both, lessens the time it takes to complete the information task. For economy of search, as measured by the number of pages accessed, spatial and/or hierarchical information seem to contribute to a more parsimonious search, that is, where users access less pages to locate target information. For both measures, the post hoc analysis of the individual groups has produced a rather counterintuitive result. Although both spatial and conceptual information is important, combining spatial and conceptual information in a single navigation tool does not produce faster or more economical search.

Regarding effectiveness measures, alphabetical information alone produces poor search accuracy relative to the spatial group. The spatial, hierarchical and spatial/hierarchical groups perform comparably on search accuracy. However, spatial information added to hierarchical information does impact users ability to recall page titles from the document. The spatial/hierarchical group retrieved significantly more titles in the free recall task than the hierarchical group.

In relation to previous research, the results of the effectiveness measures are inconsistent with McDonald and Stevenson's (1999) who found no difference on recall measures between a conceptual versus spatial map. This is possibly due to the confounding of spatial and conceptual information, which was addressed in this paper, but could also be attributed to the structural differences in the navigation aids themselves. For example, both McDonald and Stevenson's navigation tools included link designations whereas the spatial map in this study did not. Also, McDonald et al. used localized spatial maps instead of global spatial maps. Potentially, the impact of spatial information on effectiveness measures may be less crucial when the users are viewing only a portion of the whole document structure rather than the entire document structure. This is an issue that needs to be addressed in future research.

Clearly, structural aspects of a hypertext document may affect the usefulness of a given navigation tool: document size, local versus global navigation tools, the types of information provided or the linking structure. These structural effects are important to consider in the context of different navigational goals. Users begin navigating a document with an overall goal in mind. They could be browsing with no specific target in mind, searching for a particular piece of information that they have not located before, or reconstructing a previous search. The task demands of the overall goal determine the subgoals that users must set for themselves. These subgoals are not end states but are outcomes of more process-oriented strategies such as moving quickly, opening the fewest pages or remembering a page or subset of pages.

The results of the previous analyses suggest that different kinds of information in a navigation tool may support different subgoals. For example, the goal of remembering the basic elements of the document is enhanced by hierarchical information combined with spatial information over hierarchical information alone. Hence, different types of information in a navigation aid may provide support for only some of the subgoals involved in the course of navigating a document. Subsequently, the match between the goals of the user and the type of support the navigation tool was designed to provide would be particularly relevant. This would be true for both intermediate process goals as well as for overarching goals.

Given the findings of this analysis, it is apparent the appropriate question is not whether a spatial versus conceptual navigation tool is better, but rather which combinations of conceptual and spatial information provide the best support for a particular set of goals. Identifying the types of information that support intermediate processes such as moving quickly, accessing the fewest pages, recalling basic elements and accurately targeting specific information needs to be undertaken within the context of different hypertext structures across the time course of navigation.

To this end, the methodology used in this paper could be extended to study additional issues of interest such as document size, fine-grainedness of navigation tools or the integration of information and subgoals and, therefore, may be an effective approach toward understanding the complexities of hypertext use.

The implications for design highlighted by this analysis are threefold. First, providing organization based on meaningful or salient information such as conceptual or spatial information is important. Second, combining spatial and conceptual information in a single navigation tool does not guarantee improved user performance. Third, because we see differences in the information that supports different subgoals, task analysis that takes into account distinct user sub goals is important.

CHAPTER 4 – OVERT AND COVERT SPACES

4.1 The Behavioural Space

As in the approach used in the preceding chapter, much of the previous research on the effects of navigation tools looks at these effects from a performance perspective, that is, using typical measures of efficiency and effectiveness (Dee-Lucas & Larkin, 1995; Dias & Sousa, 1997; McDonald & Stevenson, 1998).

Although navigation tool studies using performance measures have yielded valuable (albeit inconclusive) information about how navigation tools impact users' ability to find specific information in a fast and economical fashion, most do not reveal much about how different navigation tools affect the paths that users choose to access target information. This is important because, although basic performance measures indicate if users' abilities to arrive at targeted information is enhanced by a navigation tool, these measures do not indicate how users went about their search or what navigation choices were made at different decision points. How does different information presented in a navigation tool affect where groups of users generally go? In order to determine this, users' path patterns must be examined. The collective path patterns for any group of users can be called the behavioural space for that group.

4.1.1 Goals of the Behavioural Analysis

In Chapter 3, it was found that spatial information does not enhance efficiency measures such as speed and economy but that, in conjunction with hierarchical information, it can enhance recall of document elements. Given that, in general, we see few differences between the spatial, hierarchical and spatial/hierarchical groups' performance, we may suppose that users' in these groups also exhibit comparable path patterns. The current chapter uses a novel method of data analysis, multidimensional scaling, to investigate if there are identifiable differences in user path patterns across the four groups.

4.1.2 Linearity of Use

Prior hypertext research has examined different aspects of users' path data with a number of approaches. Linearity of use is a focus of some of these studies. Linearity has been operationalized in a number of ways.

Beasley & Vila (1992) used proximity matrices where each cell contained the frequency of a particular page-transition, to examine the linearity of user behaviour and found that females navigated in a more linear fashion than males where linearity was defined as the choice of the next screen in a predetermined lesson sequence. It is relevant to note that, as it relates to navigation support tools, these measures of linearity did not take into account access to menu items.

Beasley & Waugh (1997) used page-transition frequencies to reveal predominant path patterns of users during learning and review. They found a strong tendency for users to navigate top-down and from left to right during initial learning but less structured navigation during review.

Horney (1993) makes the distinction between linearity of the document and linearity of user navigation. According to Horney, linearity of the document can be described through a linearity function. Path length can be derived by tracing paths backwards until the probability that a certain path is taken drops below a certain criteria.

The average length of the paths gives a measure of linearity for each node and when all nodes are averaged together, for the entire document. In contrast to this, linearity of use was based on the number of visits to each node from each of it's parents where a parent node is defined as those nodes visited immediately prior to a given node. Hence, some nodes could have more than one parent node if two nodes led into the given node. The linearity function for use was calculated from path probabilities that were derived from the ratio of parent-child traversals to the total number of node visits. Using these two definitions, Horney examined the path patterns of eight hypertext authors and found that users with different skills and goals use distinctly different navigation patterns and that these patterns show little resemblance to the linearity of the document itself.

Andris (1996) also distinguished between the linearity of the document and the linearity of the user's path. He described the baseline linearity of the document as a ratio of the minimum number of linear accesses to ensure full coverage of content vs. the minimum number of nodes accessed to ensure full coverage of both content and menu. The actual linearity of a given user's path was then defined as the ratio of actual number of linear accesses of content vs. the actual accesses of content and menu. Finally, the linearity of the user's path was defined as the ratio of baseline to actual linearity. Based on this definition, Andris found that the different learning styles of users resulted in different degrees of linearity of use.

Given that linearity of use is an aspect of interest within total document navigation, it seems prudent to examine this issue in regards to navigation tools as well. Therefore, the purpose of the present analysis is threefold: 1) to investigate if there are differences during an information search task in the page-transition patterns of users who are exposed to different navigation tools, 2) to examine if page-transition patterns reveal linear use of the navigation tool, 3) to develop an alternate approach to data analysis to help shed light on the properties that guide navigation tool use.

4.1.3 A Qualitative Approach: Multidimensional Scaling

One obstacle to such an investigation is the difficulty of making path comparisons between groups of users rather than individual users. To address this issue, a qualitative approach was taken toward examining the behavioural space of multiple users. The use of multidimensional scaling (MDS) will provide a means of making global comparisons between groups of users as well as reveal how the form of the navigation tool impacts the behavioural space of a given group.

MDS is a statistical method for revealing the underlying properties embedded within a matrix of data and displaying those properties in a visual format. In this case, the matrices of frequencies show how frequently a move was made between column and row elements. Based on the correlations between items in the frequency matrices, spatial coordinates are calculated and plotted in a configuration. MDS is often used to analyze frequency, similarity or distance data (for examples see Borg & Groenen, 1997). MDS is related to factor analyses but is more parsimonious as the optimal multidimensional solution is constrained to five dimensions or less. Solutions of more than five dimensions create spatial configurations that are not visually interpretable.

MDS is particularly suited to the frequency data of hypertext navigation because it allows us to produce a configuration that visually represents each group's path data and to partition these configurations in different ways to determine if the behavioural space of users reflects the navigation tool they used.

MDS uses measures of proximity to produce representative configurations of a data set. When we apply MDS analyses to psychological data, it is based on the assumption that measures such as similarity judgments between object pairs are equivalent to measures of proximity; that the rating or ranking of objects constitutes the psychological distance between two objects. This approach has been used in many studies of psychological phenomenon. For example, in the case of frequency data specifically, Dawson and Harshman (1986) used a variant of MDS to reveal information about letter-perception processes. In this instance, a confusion matrix contained frequencies of how often subjects confused one letter-like symbol with another. The frequencies within each cell of the matrix were indicators of the subjects' perceptions of the similarity between two symbols. The subsequent MDS analysis on this data revealed that global characteristics such as the size and general shape of a letter guided these perceptions.

In the following analysis, the frequency count is a page-transition count such as that used by Beasley and Vila (1992) and Beasley and Waugh (1997), that is, the number of times a move from one particular page to another was taken. For example, for the Alphabetical group, the average number of times subjects moved from the page entitled Club Fungi to the page entitled Conjugate Fungi was 1.73. This captures a view of the data based on transitions as opposed to states. By counting the number of particular transitions we can gain a sense of which connections are important and which are not.

Each cell of the transition matrix contains the average frequency of moves from one particular page to another. Transitions that occur frequently can be assumed to be between pages that are very related on some, as of yet, unidentified dimension whereas transitions that occur infrequently can be assumed to be between pages that are less related. In this way, the frequencies of transitions are used as a measure of proximity or relatedness.

Potentially, extraneous variables such as familiarity or response biases may impact transitions. However, if dimensions representing conceptual relationships remain consistently evident in the data it would suggest that such extraneous variables do not exert an overly strong influence relative to the importance of relatedness. Also, because transitions are averaged across subjects in a given group, such individual differences would have reduced impact.

Using this approach to analyzing frequency data, questions of interest are: Does the behavioural space reflect some aspect of the navigational tools? Is this space based on different properties or dimensions for different navigation tools? If so, what does this reveal about how different navigation tools are used? Is linearity an important aspect of navigation tool use?

4.1.4 Results

For each subject in all four navigation tool groups, a matrix was constructed of the frequencies for a move from one particular page to another. Subject matrices were then averaged to create a group matrix of averaged frequencies for each of the four navigation tool groups.

Initially, an analysis of variance was conducted on the frequency matrices for the four groups which revealed a main effect of group, F(3, 1936) = 4.053, $p <.01$. Tukey's post hoc tests (HSD) indicated this result was attributable to a difference between the Alphabetic group and the Hierarchical group, Alphabetical vs. Hierarchical, $p < .01$.

The averaged frequency matrices were then transformed into correlation matrices and analyzed using MDS.

The optimal solutions (number of dimensions) for each group were determined using several criteria typically used in MDS analyses (Borg & Groenen, 1997). The most common measure that is used to evaluate how well a particular configuration reproduces the observed distance matrix is the stress measure. Stress is a goodness-of-fit measure for the entire MDS representation based on the error between the observed proximities and the reproduced proximities that comprise the configuration. The stress criteria adopted in this study was based on the stress norms proposed by Kruskal and Wish (1978), who suggested a value of .20 or less was acceptable to define the optimal number of dimensions. Scree plots of the stress versus number of dimensions were also examined. Viewing scree plots, the optimal number of dimensions is suggested by the sharpest elbow in the curve which represents the solution where further increases in dimensions does not improve the stress statistic (Kruskal, 1964). Changes in R^2 values are also used to evaluate the precision of solutions. The optimal solution typically results in a more drastic increase in the R^2 value from the previous dimension than between any other two dimensions. For instances where two solutions were very close, Shepard diagrams, which are scatter plots that plot the proximities on the x-axis and the corresponding MDS distances on the y-axis, were also examined. The solutions that produced diagrams with the least scatter and no obvious outliers were chosen as the optimal solution. These criteria were judged collectively to determine the appropriate number of dimensions for the MDS solution.

Alphabetical: The Alphabetical tool consisted of the twenty-two page titles in a vertically organized alphabetical list (see Figure 2.4a). The list contained a beginning structure where the top label was adjacent to only one other label below it, a middle structure where labels had two adjacent neighbours and an end structure where the last label had only one adjacent label above it. No conceptual information was contained in this tool. Thus, we would expect that if users' behaviours are driven by the navigation aid, the MDS configuration should reveal the alphabetic structure. Pages that are alphabetically related should be clustered together. For example, the pages at the beginning of the alphabetical list should appear close together in the configuration, the pages in the middle of the list should be clustered and the pages at the end of the list should appear close to one another. In fact, this is exactly the type of arrangement reflected in the configurations.

According to the criteria above, the optimal solution for the Alphabetical group data was a three-dimensional solution that accounted for 75% of the variance with an acceptable stress value, Stress = .156. A scree plot of the stress values and the dimensions shows a sharp elbow at the three-dimensional solution suggesting this is the appropriate solution. The MDS configuration for the three-dimensional solution appears in Figure 7 and can be partitioned into three distinct ordered regions (for further information on regional partitioning see Borg & Groenen, 1997).

These regions represent the Alphabetical listing divided approximately into thirds in an ordered fashion; beginning, middle and end (from right to left). Although the MDS plot reflects the general form of the navigation tool in an ordered way, the order of individual points within the three regions do not precisely follow the alphabetic sequence of the list.

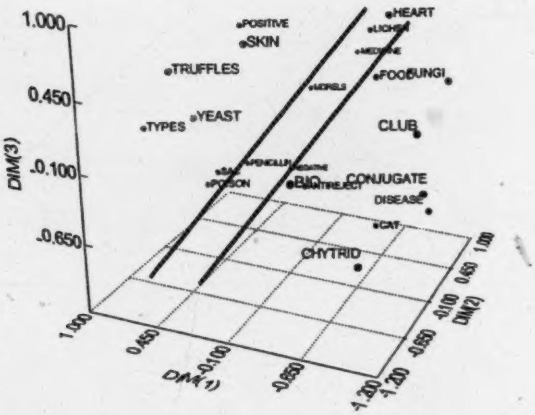

Figure 4.1. Three-dimensional solution for the Alphabetical Group – Axial Partitioning based on the divisions of the alphabetical list into thirds

In order to determine if the precise alphabetic order was possibly obscured by the three-dimensional representation, each facet of the three-dimensional plot was projected onto a two-dimensional configuration (facet diagram). Out of these 3 diagrams, the plot of Dimension 1 and Dimension 2 shows an obvious alphabetical organization into clusters of the thirds, ordered in a counterclockwise fashion (see Figure 4.2a). The plot of Dimension 1 and Dimension 3 (Figure 4.2b) shows a very clear axial partitioning of the space into the three alphabetical portions of the list. This facet is an ordered facet as well, the first third of the alphabet appearing on the right, the second third in the middle and the last third on the left. However, neither of these plots shows absolute correspondence to the order of labels on the list. The final plot of Dimension 2 and 3 shows no obvious organization based on the dimension of alphabetical order. To investigate if the third facet diagram was possibly organized via another dimension, I attempted to partition the behavioural space according to the frequency of a given page being accessed, rather than the transition frequency. Frequency of access for individual pages did not appear as a clear dimension in any of the facet diagrams for the Alphabetical group.

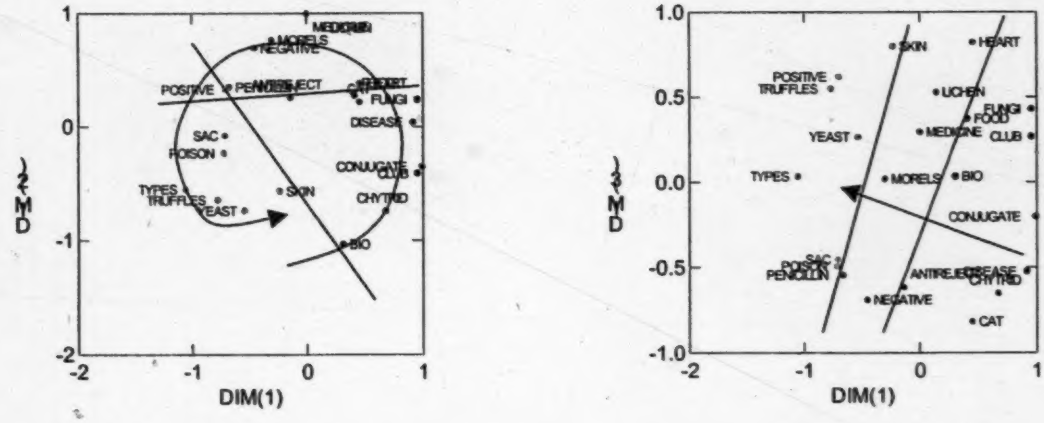

(a)

(b)

Figure 4.2. Alphabetical Facet Diagrams: (a) Dimension 1 and 2, (b) Dimension 1 and 3

Hierarchical: The Hierarchical tool was a list with the levels of the hierarchy denoted by different font sizes and label colors. This tool contained conceptual information with a minimum of spatial information. The three basic subheadings that are most salient in the list are Positive Uses, Types and Negative Aspects. If users' navigation choices are influenced by this hierarchical organization then we should expect that these three categories would be represented in the MDS configuration as well.

The optimal solution for the Hierarchical group was a two-dimensional solution, which accounted for 74% of the variance in the data. The stress value for this solution was adequate for defining an optimal solution (stress = .20). The scree plot showed a distinct elbow at the two-dimensional solution. A plot of the two–dimensional solution showed that three distinct regions could be axially partitioned (Figure 9). These three regions represent the three subheadings of the 2nd level of the hierarchy (Positive uses, Types and Negative aspects) grouped with the fourth and fifth level labels below that correspond to each. This reflects the depth of the hierarchy of the material. Partitioning of the behavioural space based on breadth (the horizontal levels of the hierarchy) did not produce any clear representations. This is not surprising as the information based on the breadth of the hierarchy was embedded in the list, denoted by font size and label color, and, therefore, was not as salient as the depth information. Frequency of access for individual pages was not reflected in the behavioural space either.

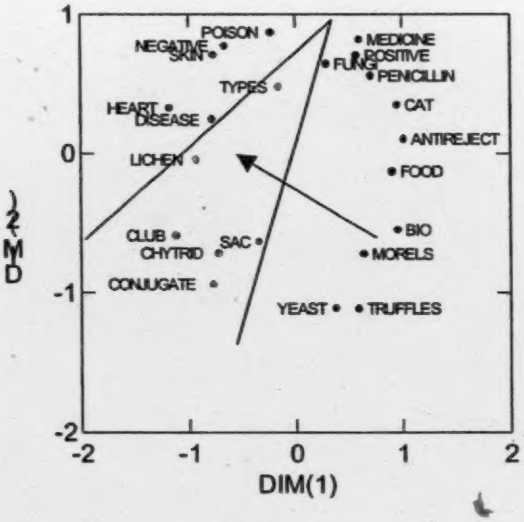

Figure 4.3. Two-dimensional solution for the Hierarchical group axially partitioned by depth of hierarchy.

Spatial: The Spatial tool was a map with no explicit links and no information about conceptual relatedness except for that suggested by the general clustering of the labels. The navigation tool contained three obvious clusters, each containing the page labels from one of the three main categories. If users' decisions are affected by these spatial groupings, we should see these clusters reiterated in the MDS configuration and this is exactly what occurs.

The best solution for the Spatial group is also a two-dimensional solution, $R^2 = .870$, stress = .183, with a scree plot elbow at the two-dimensional mark. The two-dimensional plot for the Spatial group showed similar groupings to the hierarchical group and could be partitioned in two different ways, axially and in clusters. The axial partitioning of the space clearly delineates the three depth subheadings of the hierarchy (see Figure 10) but cannot be partitioned clearly into the horizontal levels of the hierarchy. The clustering of regions also reflects the vertical levels, although the cluster regions do not take into account the FUNGI label. These clusters closely resemble the clustering of the page titles on the spatial navigation tool but the individual points do not adhere to the vertical order of the labels.

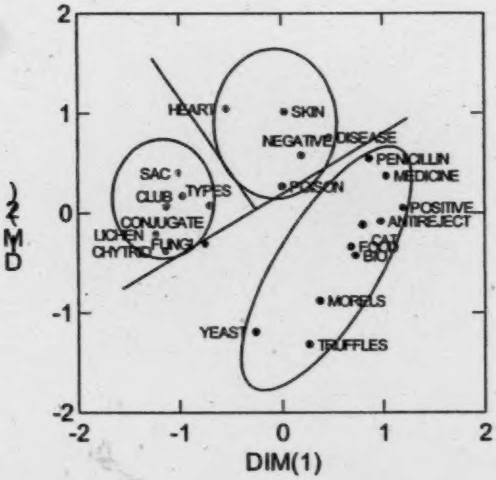

Figure 4.4. Spatial two-dimensional solution: Axial partitioning and clustering according to three depth levels of hierarchy.

Spatial/hierarchical: The spatial/hierarchical tool was a tree-diagram that explicitly showed the hierarchical categories and their relations to one another. This tool contained both conceptual and spatial information. The conceptual information is organized in an explicit hierarchy, making available to the user information about both the depth and breadth of the organization whereas the hierarchical tool conveys only information about the depth of the hierarchy. Therefore, we should expect that the behavioural space for the spatial/hierarchical group could be partitioned into areas that independently represent depth and breadth.

The best solution for this group is a three-dimensional one, $R^2 = .83$, stress = .135 with a scree plot elbow at the three-dimensional mark. The three-dimensional plot (see Figure 11) could be partitioned both axially and radially into the three subheadings of the hierarchy.

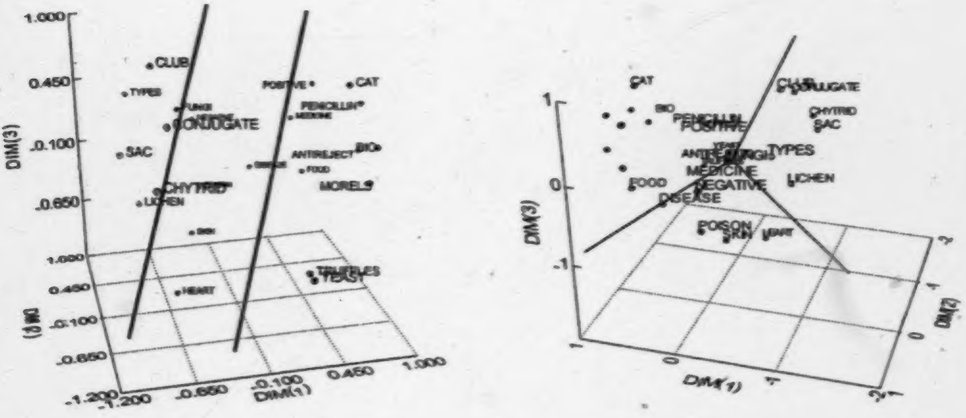

a)

b)

Figure 4.5. Three-dimensional MDS solution for the Spatialhier group – (a) Axial partitioning based on depth of hierarchy, (b) Radial partitioning by depth from Fungi center point.

This type of partitioning illustrates users' reliance on information provided by the depth of the material. The dimension of breadth is also discernible in the three-dimensional plots for this group as illustrated in Figure 12, which shows a distinct partitioning of the behavioural space into the three horizontal layers of the hierarchy. However, this partitioning is not absolute as the point "Biocontrol" is outside the appropriate region, yet inspection of the "Biocontrol" label on the spatial/hierarchical tool provides no obvious reason for this.

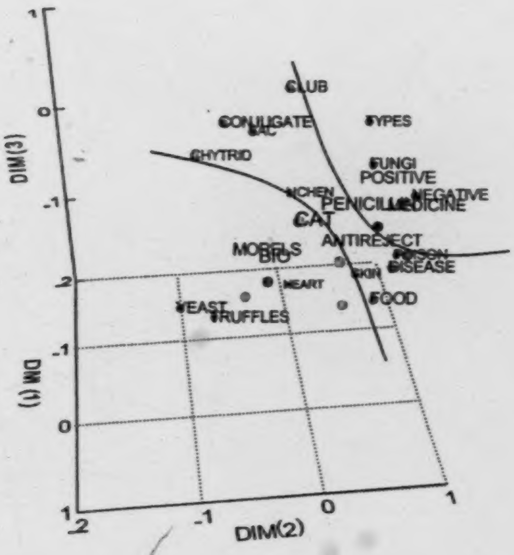

Figure 4.6. Three dimensional MDS solution for the Spatialhier group - Axial partitioning by breadth of the hierarchy

In the case of the spatial/hierarchical tool, both depth and breadth are explicitly presented in the navigation tool and both types of information are also evident in the users' transition choices.

Linearity: Because the page-transition frequencies so obviously mirror the navigation tool structures, it is tempting to deduce that users are merely moving through the navigation tools in a linear fashion. However, a purely linear usage of the navigation tool would result in the one-dimensional MDS solution being the most optimal for describing the data. This would be particularly true for the two lists which users are limited to traversing vertically. Consequently, if there is a significant difference between the one-dimensional solution and the optimal solution chosen for each group, this would suggest that another property, potentially in conjunction with linearity, was guiding tool use. To ascertain if users moved solely in a linear manner, F tests were conducted on the variance for the one-dimensional solution versus the variance for the optimal solution for each group. For all four groups a significant difference was detected between the R^2 values, (Alphabetical, $F(3,18) = 12.05$, $p < .05$, Hierarchical, $F(2,19) = 24$, $p < .05$, Spatial, $F(2,19) = 25.14$, $p < .05$, and Spatial/hierarchical, $F(3,18) = 5.22$, $p < .05$). Therefore, linearity is not the sole dimension that influences how users traverse the navigation tool.

4.1.5 Discussion

The structure of the navigation tool heavily influences the transition choices that users make as indicated by the relationships found between the spatial configurations derived from the page-transition frequencies and the organization of the navigation tools. The MDS configurations for each group can be clearly partitioned into dimensions that reflect the elements contained in each of the four navigation tools demonstrating that different navigation paths were produced by presenting different navigation tools to users. Even though users were unconstrained by the linkage structure of the pages, they still relied heavily on the structure of the tool to guide their choices. This was true even for the groups that used the two lists and the spatial tool, which did not contain any explicit linkage suggestions within the tool. With these tools, there were no elements to suggest some pages were linked while others were not. Hence, users were free to access any page title without any suggestion that certain pages were not directly accessible from other pages. This allows them the freedom to make choices based solely on the potential content of a page rather than on the link connections between pages. However, the results show that users did not exploit this freedom but rather were guided substantially by the order and structure of the navigation tool.

Looking closely at the MDS solutions, there is a difference in the number of dimensions for an optimal solution in the alphabetical and spatial/hierarchical groups versus the hierarchical and spatial groups. This difference suggests that users' are making different transition choices between the four navigation tools. The optimal solution for the Alphabetical tool group was a three dimensional solution where the most identifiable dimension was alphabetic order. For the hierarchical group, a two-dimensional solution produced the clearest organization of the data, the important dimension being depth of the hierarchy. The spatial group data was best represented by a two-dimensional solution as well where the depth of the hierarchy was clearly identifiable as a guiding dimension. Finally, the optimal solution for the spatial/hierarchical group was three-dimensional with the depth and breadth of the hierarchy both represented in the facet diagrams. For this last group, if the breadth representation was just a matter of users accessing the lowest level of the hierarchy infrequently relative to the other levels, this would also be reflected in the plots partitioned by the frequency of access for individual pages for the other navigation tools. This is not the case. Therefore, this suggests that, for the spatial/hierarchical group, users rely on the breadth information in the tool to make navigation decisions as well as the depth information.

Generally, users' navigation transitions seem to reflect the most prevalent organization in a given navigation tool whether that is conceptual, alphabetical or spatial. However, there are two different lines of evidence to suggest this is not totally linear usage of the tool. First, if use of the navigation tool were absolutely linear then the MDS analyses would reveal a one-dimensional solution as optimal; this is not the case for any of the groups. Secondly, inspection of the plots of the optimal solution for each group should reveal not only partitioning based on salient features of the navigation tool (e.g., the three depth subheadings) but also all of the individual points lying within certain regions should precisely correspond to the order of the labels in the navigation tool as well. Again, this is not the case. For example, with the Spatial tool, linear navigation would result in the "column" on the left side of the tool being traversed from top to bottom, followed by a move to the right to the top of the next "column" (see Figure 2.5). The order of navigation would then be Caterpillar, Penicillin, Anti-rejection Drug, Medicine, Biocontrol, Positive Uses, Food, Morels, Truffles, Yeast. However, the two-dimensional-solution plot of where users actually went (Figure 4.4), illustrates that, in the cluster that contains these particular labels, the above order of navigation is not represented. These two lines of evidence indicate that, although the property of linearity may contribute to navigation, it does not function in isolation but rather in conjunction with other guiding dimensions such as the conceptual layout of the material or the user's individual mental representation of the space.

Summarily, information contained within the structure of the navigation tool and linearity of use explains much of the observed navigation behaviour. A remaining question is, does the inherent organization of the content affect navigation, that is, are users relying on content to make any navigation decisions? The data from the alphabetic group provides an opportunity to speculate on this issue. Besides random choices, users of the alphabetic navigation tool are limited to reliance on three sources of information for navigation choices: 1) alphabetic order, 2) linear order or 3) content as suggested by the page titles. If the alphabetic group was navigating only linearly or only alphabetically, the best MDS solution would be one-dimensional. Because this is not the case, a conceivable explanation is that users must be making some transition choices based on content as well, that is, a best guess from the labels or previously visited pages as to what would be an appropriate page to find the target information.

In summary, three conclusions are suggested by the results of this configural analysis of page-transition data. First, users' transition choices are heavily influenced by the types of information presented in the navigation tool. Second, users' do not adhere strictly to a linear usage of the navigation tool. Third, in an unfamiliar document, even with a navigation tool that presents limited content information (only the page titles themselves), users seem to make some page-transition choices based on suppositions of content.

These results may seem somewhat obvious but the MDS analysis gives reason to suppose that this relationship may be more complex than expected. For example, the optimal solutions for each of the four groups are of different dimensions, in some instances two dimensions describe the data, in other instances three dimensions are required. If navigation behaviour was wholly dependent on the tool, the optimal MDS solution to describe that behaviour would be two dimensions for each of the groups, in essence, a reproduction of the two-dimensional configuration of each navigation tool. These are important findings because they suggest that other properties or characteristics of the course of navigation are mediating the relationship between the navigation tool and the resultant behaviours. Future research should focus on identifying those variables and understanding their role in the sequence of relationships between navigation tool, behaviour and mental representations.

This analysis also pioneers the use of MDS on page-transition data as a viable and enriching approach toward investigating the properties that guide navigation tool use. In this study, MDS analysis of page-transition data revealed highly systematic and interpretable regularities in the flow of user transitions. This type of scrutiny is not possible using more traditional methods of group comparison. Other commonly used statistical methods, like analysis of variance, fail to capture all the differences between groups of users' page-transition patterns, whereas MDS represents the data in a configural manner that highlights these differences and allows for the detection of theoretically relevant properties that underlie the data.

4.2 The Representational Space

Chapter 3 provided evidence that the information in a navigation tool impacts some aspects of performance. From the results reported above it is apparent that the actual path patterns of users, that is, the behavioural space of a given group, can be influenced by the form of the navigation tool as well. These findings lead to an interesting question: Are navigation tools impacting navigation behaviours because they influence users' underlying mental representations of the document or do all users develop the same representation but exhibit different behaviours?

Several studies suggest links between graphical tools and users' abilities to recreate the structure of a hypertext document and also differences in performance between users who develop graphical representations and those who don't. McDonald and Stevenson (1997) found that users given a spatial map to help with an information search task subsequently placed more correct node titles on a map of the document than users given a content list or no navigation tool. Kerr (1990) found that faster users have more accurate and more graphically detailed representations, while slower users tended to give verbal descriptions of the document. These studies imply that navigation tools can facilitate the use or development of accurate mental representations of a hypertext document.

However, other studies provide contrary evidence. Stanton, Taylor and Tweedie (1992) found that users given a map for navigation were less able to recreate a map of the system than users who were given no navigation aid. Specifically, map users drew fewer primary links and total number of links on an outline of the hypertext system than users without maps. Wenger and Payne (1994) found that providing a graphical browser for users did not improve their recall of text structure over that of users given no navigation aid.

Previous methods for accessing users' conceptualizations of the document have included users creating sketch maps (Gray, 1990), filling in blank tree diagrams or outlines of the document (McDonald & Stevenson, 1997), arranging page labels in a configuration, giving verbal descriptions of the format of the document (Kerr, 1990). Such methods tended to force users to create a spatial configuration of the document whether or not the navigation tool used was spatial in nature.

In this study, as described in the methods chapter, users were asked to complete a ratings task which required them to rate the distance between 100 pairs of page titles (e.g., Lichen to Conjugation Fungi). The ratings task was administered with two sets of instructions. These instructions were designed to allow the user to interpret the task as either a spatially based or a semantically based task. Users were asked to rate the distance between pairs of pages based on either their closeness (spatial relation) or relatedness (semantic relation). The premise of this task is that users' ratings of the distances between pages will reflect the underlying representational space they have formed of the document.

4.2.1 Goals of the Representational Analysis

The goal of this analysis was to answer two specific questions : 1) Are the ratings results different between groups? 2) is there a relationship between the representational space (ratings results) and the behavioural space (actual path patterns)?

Unlike the behavioural space analysis, the data for the representational space was not conducive to the use of MDS as the ratings task only produced data for a portion of all possible combinations of pages. This resulted in a data matrix with a large amount of empty cells (a sparse matrix), making the use of MDS unreliable. Instead, the ratings data was examined using correlations and regression analyses.

4.2.3 Results

Out of the 100 pairs, four pairs were omitted due to repetition or missing labels. Correlations for the average ratings between groups were high, indicating the patterns of ratings were similar across groups (see Table 4.1).

Table 4.1.

Correlations for average ratings between all groups.

	Alpha	Hierarchical	Spatial	Spatial/ Hierarchical
Alpha	1.000			
Hierarchical	0.864	1.000		
Spatial	0.802	0.853	1.000	
Spatial/hierarchical	0.707	0.819	.877	1.000

Instruction conditions

For each of the four groups, two-sample t-tests were conducted to detect any differences in ratings between the two instruction subgroups. Three of the four groups showed no differences between question types. Only the hierarchical group showed a significant difference between the question types, $t(190) = -2.55$, $p < .05$. Why might we see a difference in this group? Possibly, the conceptual information in the hierarchical aid, that is, the hierarchical organization of subheadings, was salient enough to the users that a distinction could be made between the notion of "closeness" versus the notion of "relatedness". With the spatial tools, for the "related" condition, it is possible the users assumed the spatial layout was a representation of the conceptual relatedness of the pages and, therefore, made no distinction between the terms "close" and "related". With the alphabetical list, either term might appear ambiguous.

The Relationship between Behaviour and Representation

The MDS analyses on users' path patterns at the beginning of this chapter revealed that users' sequence of movements closely reflected the structure of the navigation tool. Given that the information in the navigation tool influenced where users went, the relationship between where they went (frequency of page-transition) and how they perceive pages to be related (ratings task) should be suggestive of the type of information that is the basis for their underlying mental representation. Multiple regression analyses were done to determine if frequencies (behaviour) predicted ratings (mental representation). The analyses reveal that for the two groups that contained spatial information, the spatial and spatial/hierarchical group, the page-transition frequencies predicted the users ratings of distance between pages, Spatial, $R^2 = .387$, $F(1, 94) = 59.438$, $p < .001$, Spatial/hierarchical, $R^2 = .141$, $F(1, 94) = 15.373$, $p < .001$.This suggests that the experience of moving through the document contributed to the users' mental representation of the document space. That is, their perception of the document structure was at least partially motivated by their behaviour within the document, which was in turn influenced by the navigation tool they used. This was not true for the Alphabetical and Hierarchical groups. One possible account for this result is that, spatial information may facilitate the integration of knowledge obtained from two separate sources: the navigation tool itself and the experience of moving through the document. If so, users given spatial information have the advantage of only having to access a single representation in memory rather than having to mentally combine information from multiple sources.

4.2.3 Discussion

Based on the distance ratings, high correlations between the four groups' average ratings indicate the patterns of ratings were comparable across groups suggesting all groups formed a similar mental representation.

At the start of the tasks, the users are first exposed to the navigation tool. The assumption is that, initially, the navigation tool prompts users to form a hypothetical organization for the material. As users move through the document, their navigation experiences potentially confirm and facilitate further development of these representations. Although the mental representations as indicated by the ratings (the representational space) are similar across groups, the page-transition frequencies (the behavioural space) indicate that navigation behaviours are differentially associated with the structure of the navigation tool. For spatial tools (spatial and spatial/hierarchical), the representational space and the behavioural space are related. Users' navigation behaviours predict the mental representations they have formed by the completion of the tasks. For the non-spatial tools (alphabetical and hierarchical), the representational space and the behavioural space do not appear to be associated. Users' navigation behaviours do not predict the mental representations they have formed by the completion of the tasks.

4.3 Summary and Implications

To summarize the results of the behavioural data collected thus far: First, navigation tools containing spatial and hierarchical information enhance performance over an alphabetically ordered tool whether these types of information appear in isolation or combination. However, combining spatial and hierarchical information does not produce enhanced performance over either alone.

Second, as evidenced by the MDS analyses, where users go is substantially influenced by the information in the navigation tool. The behavioural space of a given group reflects the predominant type of information within that particular navigation tool.

Third, even though users in all four groups produced comparable ratings on the ratings task, only the two groups given spatial information exhibited navigation behaviours that predicted their subsequent mental representations of the document.

4.3.1 Performance and Representation

If all groups form comparable mental representations of the document material, why would the alphabetical group perform poorly relative to the other three groups? A possible explanation is that, for the alphabetical and hierarchical groups, two representations develop, a representation of the navigation tool itself and a representation of the content of the document. For the hierarchical list users, the presence of hierarchical information provides conceptual information related to the information conveyed by the two spatial tools, possibly allowing for some merging of the two representations. Hierarchical organization could be suggestive of a tree-like organization with superordinate, basic and subordinate levels. Hence, the hierarchical group performs comparably to the two spatial groups. For alphabetical users, the absence of information that organizes the material in a meaningful way leaves the user with separate unrelated notions of the document space, one based on the alphabetical organization of the content list and the other based on a relational organization constructed through the experience of navigation.

If this is the case, then discrepancies in performance, as reported in Chapter 3, would not be due to the development of different mental representations but may instead be due to difficulties in consolidating multiple, unrelated forms of information (e.g., the alphabetical list has no relation to the semantic-relatedness of pages) into a single, cohesive representation. This suggests that performance is enhanced when the information in the navigation tool corresponds to the conceptual organization of the material and that conceptual relationships can be conveyed via spatial or hierarchical information.

4.3.2 Behaviour and Representation

The regression analyses on the ratings data indicate that behaviour predicts mental representation for the two groups that contain spatial information. The navigation tool with hierarchical information alone does not produce this relationship nor does the alphabetical tool. This implies that, in the ongoing links between the navigation tool, behaviour and representation there is something unique about spatial information. It is possible that spatial tools provide an initial hypothetical organization of the material in a form that facilitates the integration of further information, that is, relational information gathered through traversal of the document. This results in a final integrated representation of information from all three sources.

The analysis of the representational data backs up the complexity of the relationships between different aspects of the navigation experience as suggested by the MDS behavioural analysis. There appears to be a continuous relationship between elements of the early navigation experience and variables that appear later on. However, this does not seem to be a one-to-one correspondence between individual variables but rather a cumulative, sequential relationship between numerous variables.

Although I have provided a possible explanation for the discrepancies between behaviour and representation, this explanation is speculative. Also, the ratings task rests on the assumption that peoples' estimates of distance between pages reflects their actual underlying representation. Questions regarding the mental representations of users cannot be satisfied from this data alone. In order to better understand how such discrepancies can occur this issue should be viewed from a broader theoretical perspective. This is the theme of the following chapter.

CHAPTER 5: REPRESENTATION AND BEHAVIOUR: A SYNTHETIC APPROACH TO AN EARLY DEBATE

5.1 Introduction

Chapter four provides behavioural evidence that suggests the links between behaviour and representation in a hypertext environment are somehow mediated by spatial information. However, users provided with various types of information seem to formulate comparable representations to those given spatial information even though their navigation behaviours are different. The obstacle in establishing explanations for such behavioural/representational discrepancies is that, of course, the cognitive system we are speaking about is the human brain and we have no direct means for viewing mental representation in the human brain. On the other hand, a computational framework such as a connectionist system provides an opportunity to examine underlying representations within an information processing system. Consequently, the current chapter presents an alternate methodology for investigating discrepancies between representation and behaviour. It focuses on spatial representations in particular for two reasons: 1) the predominance of the spatial metaphor in hypertext literature brings up questions regarding the definition of space and the processing of spatial information, 2) as evidenced by the ratings data in Chapter four, spatial information seems to play a distinctive, albeit unexplained, role in hypertext navigation.

As outlined in chapter one, there are several different ways of conceptualizing the properties of space and the course of spatial knowledge acquisition. One such supposition is that spatial encoding involves internalizing the metric properties of an environment and that the resulting mental representation of that environment possesses metric properties as well. Such proposals are not without controversy. Nonetheless, the notion of metric spatial representation has been studied in several related contexts, some of which were superficially discussed in the introduction of this thesis. The following sections review in more detail the nature of these representations in three key areas of study: similarity spaces, mental imagery and cognitive mapping.

5.2 Similarity Space Representations

The previous chapter investigated the behavioural space of users under the assumption that the frequency of their page-transitions reflected the perceived relatedness of the two pages traversed. The subsequent MDS analysis of this data is based on the premise that the psychological distance or similarity between entities can be represented as points in a multidimensional space. Points that represent related concepts are close together in the similarity space, whereas points that represent very different concepts are far apart. This similarity space is assumed to have metric properties. Specifically, this space adheres to the three properties of minimality, symmetry and triangle inequality. The concept of a similarity space has been linked to explanations of many types of cognitive phenomena such as similarity judgments (Tversky,1977, 1982, 1986; Krumhansl, 1978, 1982), analogical reasoning (Sternberg, 1977) and judgments of metaphor aptness (Tourangeau & Sternberg,1981, 1982).

An assumption that emerged from the similarity space proposal was that the behaviour based on this space would mirror the metric properties inherent in the underlying representational space. More simply, similarity space mediates similarity judgments. Behavioural evidence for and against the similarity space proposal was generated through studies of human judgments of similarity.

Tversky (1977,1982,1986) found there were a number of variables that could consistently impact people's similarity judgments, resulting in responses that violated the principles of metric space. For example, with respect to the symmetry principle, an unfamiliar category is judged to be more similar to a familiar category than vice versa. Tversky argued that such behavioural discrepancies were evidence that the mental representations that supported similarity judgments could not be metric.

In response to Tversky's assertions, Krumhansl (1978,1982) provided evidence indicating that if additional information is taken into account when evaluating similarity judgments, judgments that are nonmetric can still be based on a metric space. Specifically, if the relationship between points in the similarity space is evaluated not only by distance but also using procedures that are sensitive to density, how crowded the space is in different positions, then the metric space can lead to similarity judgments that are not symmetric.

The similarity space proposal was based on the human experience of similarity, an experience that is not distinctly spatially based. The concept that mental representations of space could be organized in precisely the same way that physical spaces are organized was even more apparent in the study of mental images.

5.3 Mental imagery

The theme of disparities between representation and behaviour was echoed in another early debate centred around the nature of mental images.

Kosslyn, Ball and Reiser (1978) produced behavioural evidence that suggested mental images of spatial configurations preserved the metric qualities observed in real-space. For example, images seemed to have the property of distance because subjects took longer to scan across larger distances in a mental image than shorter distances. Images also seemed to have the property of size as it took subjects longer to report small details of small images than small details in a large image. Shepard (1962) found that subjects took longer to mentally rotate one object to match another when the disparity of the degree of rotation between them was large.

Counter to Kosslyn's position, Pylyshyn (Pylyshyn, 1980, 1981, 1984, in press) argues that mental images are merely "mental structures that play a role in an information processing account of mental activity" (in press) and that the actual representations of knowledge are of another form. He emphasizes the distinction between the "functional space", that is, whatever representation (e.g., propositions, Cartesian coordinates or a matrix of distances) carries the crucial descriptors of a space, and the two-dimensional pictorial expression of those descriptors. The pictorial space (mental image) is bound by the physical laws that are inherent in the environment represented. Contrary to this, the functional space is not bound by physical laws. For example, places in the pictorial space can be between other places within that space, which necessitates travelling through one place to access another. In the functional space, such as a distance matrix, there are no such constraints. From this perspective, there is no reason to suppose that observed behaviour must reflect the properties of the underlying representation.

The pursuit of a resolution to the imagery debate prompted some researchers to consider evidence gathered from the area of cognitive neuroscience. Current brain imaging techniques have shown that mental imagery elicits activity in many of the brain areas involved in visual perception, in particular the primary visual cortex. These areas of the visual cortex are spatially organized, both physically and functionally. (Farah, 1989; Kosslyn, 1994, 1995, 1997, 1999; Thompson, 2001). Kossyln relies on such evidence to support his proposal that mental images are patterns of activity in a spatially organized structure in the occipital lobe. However, findings from neuroscience studies are not conclusive as several researchers have provided evidence that these same areas can be activated by other cognitive processes and not all mental imagery tasks elicit activity in the primary visual cortex (for review see Mellet, 1998).

The previous literature on similarity space and mental imagery illustrates the potential for discrepancies between the form of spatial mental representation and the behaviour it produces. How are these notions of mental representation related to actual navigation behaviour?

5.4 Cognitive Mapping

The idea that intelligent agents develop spatial mental representations that preserve metric properties was initiated by Tolman's (1948) cognitive map proposal. Based on his studies of rat behaviour, Tolman concluded that some type of mental map that mirrored properties of the environment was established in the brain during the course of learning. This map allowed the animal to generate new navigation behaviours that were not based on the animal's direct experience of two places. For example, animals were able to generate novel short cuts between places to reach desired locations. Many subsequent studies of animal behaviour demonstrate that animal representations of space do indeed preserve metric properties (for introductions, see Cheng & Spetch, 1998; Gallistel, 1990).

Consequently, neuroscientists have been interested in identifying the brain mechanisms that may be responsible for encoding metric information. In animal studies, recordings of single-cell activity in the hippocampus suggest this structure may be responsible for the physical realization of the metric cognitive map (for review of hippocampal cell activity see Best, White & Minai, 2001). These individual cells, referred to as "place cells", were found to be sensitive to specific configurations of distal cues in the environment and reacted in a predictable fashion to changes in those cues. Based on such findings, O'Keefe and Nadel (1978) proposed that the hippocampus is the neural foundation for cognitive maps and that place cells are the primary units of the map. Further research (O'Keefe & Burgess, 1996) revealed that the receptive field of a place cell can be described as the sum of two or more Gaussian tuning curves sensitive to the distance of the animal away from a wall in its environment. Hence, the environment is represented in the hippocampus via collections of place cells, each of which represents a specific spatial region of that environment.

Regarding theories of human navigation, in early theories (e.g., Siegel & White, 1975, Hart & Moore, 1973) the role of metric information in knowledge acquisition of large-scale environments was considered the qualitative outcome of the accumulation of other types of spatial knowledge. These theories rested on the supposition that spatial knowledge is acquired in a sequence of qualitative transitions. Initially, individual places or landmarks are encoded. These are discrete, recognizable patterns of objects in the environment that, in and of themselves, do not contain information about the complete space. Beyond landmark recognition, sequences of landmarks (routes) are encoded, that is, landmarks are linked together by actions that occur between them. Finally, collections of routes are encoded which represent a survey perspective of a specific environment including the metric relationships between places. This view of spatial knowledge acquisition assumes that only nonmetric knowledge exists during the earliest stages of learning. The accumulation of nonmetric knowledge eventually results in the acquisition of metric knowledge.

More recently, Montello (1993) proposed that metric information plays a more continuous role in spatial knowledge acquisition. According to Montello, metric information is encoded in a separate spatial memory system where nonmetric information is neither a precursor for nor an inherent part of metric knowledge. Montello supports his theory with evidence from a number of studies demonstrating that even with minimal exposure to an unfamiliar path, subjects are sensitive to differences in metric properties (e.g., segment length and turn angles) in two paths that are topologically identical (Klatzky et al., 1990; Montello & Pick, 1993; Sadalla & Montello, 1989; Smyth & Kennedy, 1982). Such results could not occur if nonmetric information (landmark and route knowledge) was a necessary prerequisite for metric knowledge acquisition. Rather, these studies suggest that metric knowledge is acquired concurrently with nonmetric knowledge.

This review of similarity spaces, mental imagery and cognitive mapping suggests two characteristics of spatial mental representations: 1) that metric spatial representations do not necessarily initiate metric behaviours and, 2) that metric knowledge is an essential component of spatial mental representation from the onset of exposure to a novel environment. The latter statement provides impetus for further study of metric knowledge acquisition while the former statement provides a specific line of inquiry, which is investigated through the use of computer simulations described in the next sections of this chapter.

5.5 A Synthetic Approach

If an intelligent agent retains metric information as part of the mental representation of spatial relations, does that representation automatically produce behaviour that adheres to metric principles?

To investigate this question, an alternate approach to the traditional behavioural methodology, a synthetic approach, was employed. Why is a synthetic approach warranted? The evidence produced by traditional cognitive psychology paradigms has not resolved the debate on the role of metric information in mediating spatial behaviour. These traditional methods are based on the decomposition of complex human behaviour into basic processes that suggest the form of the underlying representation (Cummins, 1983; Dawson, 1998).

Although this approach to analysis has yielded a large and informative body of knowledge, discrepancies between perspectives such as similarity, mental imagery and cognitive mapping may be due to two issues related to these established methods. First, it is difficult to attribute observed behaviours specifically to one of three possible origins: behaviour generated by the organism, behaviour prompted by the environment and behaviour that reflects an interaction of the two (Braitenberg, 1984; Simon, 1996). Second, the attribution of behaviours to a specific origin is theory-driven, therefore, there is always the potential for uniformities in the data that are counterintuitive to theory to go unnoticed. Because a synthetic approach involves constructing a theory-neutral system that allows for the examination of the processing undertaken by different components of the system, that is, the identification of the origins of its behaviour, this approach can be used to address these issues.

A simulation approach is particularly suited to the metric representation example explored in the introduction of this paper because it is possible to construct a system that directly adheres to the properties inherent in the formal definition of space, that is, metric space. These are the same properties explored in much intensive psychological research devoted to human spatial processing as described in the previous sections on similarity, mental imagery and cognitive mapping. Although it may be argued that geographical space is an arbitrary construction, with inaccuracies embedded within it, it is a widely-studied and standardized system. Physical space adheres consistently to the three axioms of minimality, symmetry and triangle inequality. This formal definition provides an opportunity to use a set of reliable statistical tools for further analysis.

Contrary to physical space, hyperspace has no standardized and widely accepted definition with specific properties. Because of this ambiguity, decisions for appropriate network representations based on the "Fungi" hyperspace could have been compromised on several levels. For example, which of the four navigation tools most accurately represents the essence of the hyperspace defined by the web-site? What measurement unit is appropriate to represent the properties or qualities of relationships between entities in that hyperspace? What is the nature of the global referent system that encompasses the relational or local space of this particular hyperspace? In this preliminary network study, the number of assumptions required to establish a local and global referent system for a particular hyperspace may be numerous. Accurate representations of external information that retain the most essential properties of a particular real-world phenomena are critical to meaningful simulations of human cognition. Network processing is only as stable as the information used as input.

5.5.1 Goals of the study

Chapter one raised several questions regarding the nature of hyperspace from a psychological perspective. A primary question is – Does hypertext navigation involve spatial processing? In order to answer this question fully, it is necessary to first understand precisely what information is being represented and manipulated during spatial processing is the goal of the following network study.

In particular, I describe two parallel-distributed processing (PDP) networks that were trained to rate the distance between pairs of places in Alberta. There were several questions of interest. First, if very few constraints are placed on the internal structure of a system, then will it develop metric representations to guide its behaviour? Second, if the same system is designed to generate behaviours that violate metric properties of space, then will it develop nonmetric internal representations? Or can both metric and nonmetric behaviours be mediated by similar (metric) representations?

5.5.2 Introduction to PDP Networks

Generally, a PDP network is a multi-layered system that consists of a layer of input units that receive information from the environment, a layer of hidden units that process the information and a set of output units which produce a response based on the activation patterns produced by the entire network. Each of these units is connected to the other units by connection links. The strength of a connection is represented by a connection weight that can be positive or negative causing a particular input to either excite or inhibit the unit that receives it (McClelland, 1989). Each processing unit in a PDP network should be viewed as a representation of a single neuron, which could potentially be connected to many other neurons, each receiving their own input from other neurons and each producing their own outputs, which are passed onward to other neurons.

Because PDP networks are computational models, the information that is introduced into and processed by the system are mathematical representations of information relevant to the task to be performed. Each input pattern is a mathematical representation of the current stimulus of the information processing task.

PDP networks are adaptive systems, that is, they begin with a set architecture (the arrangement and number of units and the initial pattern of connectivity between units) whose form and processing can be progressively altered through the presentation of information. Thus, PDP networks are capable of learning, gradually discerning over time the underlying patterns and structure within the information received and altering responses accordingly. Unlike the processing implemented in a typical computer program where the program provides explicit information on how a task should be executed, PDP networks are trained to accurately complete a given task. The behaviour of a single processing unit in this system can be characterized as follows: First, the unit computes the total signal being sent to it by other processors in the network (net input). Second, the unit adopts a particular level of internal activation on the basis of this computed signal. Third, the unit generates its own signal that is based on its level of internal activity, and sends this signal on to other processors.

A common learning algorithm for presenting feedback to the network is called the generalized delta rule (Rumelhart, 1986). It is based on the error between the network's responses and the correct responses for a given input pattern. As each pattern is processed through the network, an error term is generated that reflects the difference between the actual output computed by the network and the target response set out for that particular input pattern. The error signal is sent backward through the network from the output units to the hidden units. Each hidden unit computes its overall error by treating the incoming error signals as net input (i.e., a hidden unit's total error is the sum of the weighted error signals that it is receiving from each output unit). Once a hidden unit has computed its overall error, the network responds to the discrepancy by adjusting the connection weights between the output and the hidden units so the output is either more or less activated depending which will lessen the error. Assuming the hidden unit, not just the output unit, contributes to the overall error, the weights leading into the hidden units are adjusted as well, based on the errors of the output units and the weights between the hidden and output units. In this way, each time a particular input pattern is presented to the network the connection weights will adjust to lessen the error. By repeating this procedure a large number of times for each pattern in the training set, eventually, the error will be so negligible that we can judge the network as having learned the appropriate output response. When the network can discern all the correct responses for all the input patterns associated with a given task, we can say the network has been trained on that task. At the end of training, the network will have a very specific pattern of connectivity in comparison to its random start.

5.5.3 Value Unit Architecture

The hidden units in a PDP network use an activation function to calculate the signal to be forwarded to the output units. Most networks using the generalized delta rule for training have hidden units that use a monotonic activation function, that is, as the input increases, the activation function increases as well. For this reason, units that use this activation function have been called integration devices (Ballard, 1986).
and their activation function is a sigmoid-shaped curve that is defined by the logistic equation (Rumelhart, 1986).

However, some networks use processors that are tuned to activate to a small range of net inputs, and generate weak responses to net inputs that are either too small or too large to fall in this range. The response of such a processor is nonmonotonically related to increases in net input, and for this reason it is called a value unit (Ballard, 1986).

One example of a network of value units is an architecture designed by Dawson & Schopflocher (1992). The processing units in this architecture use a Gaussian activation function that ranges between 0 and 1, with a standard deviation of 1. These units generate a maximum response of 1 when their net input is equal to the mean of the Gaussian. Networks of value units are trained with a variation of the generalized delta rule.

Several studies have illustrated that the internal representations produced by value units networks are more easily interpreted than those of other types of networks. (Berkeley, 1995; Dawson, 1996, 1997, 1998, 2000; Leighton, 1999; Zimmerman, 1999). Thus, the value unit network is particularly well suited to address questions related to the nature of representations.

5.6 Simulation 1: Metric Spatial Judgments

5.6.1 Methods

In the first simulation described below, a network was trained to make judgments about the "crows flight" distances between places on a map, specifically, thirteen different locations in the province of Alberta: Banff, Calgary, Camrose, Drumheller, Edmonton, Fort McMurray, Grande Prairie, Jasper, Lethbridge, Lloydminster, Medicine Hat, Red Deer, and Slave Lake. All the possible pairings of these thirteen cities resulted in 169 distances. This included distance ratings between each city and itself, which would be zero. Each pair of city names appeared twice, once for each order (e.g., the distance between Banff and Calgary would be rated, as would the distance between Calgary and Banff). The actual distances between cities in kilometres are shown in Table 5.1.

	BANFF	CALGARY	CAMROSE	DRUMHELLER	EDMONTON	FT. MCMURRAY	GR. PRAIRIE	JASPER	LETHBRIDGE	LLOYDMINSTER	MEDICINE HAT	RED DEER	SLAVE LAKE
BANFF	0	128	381	263	401	840	682	287	342	626	419	253	652
CALGARY	128	0	274	138	294	733	720	412	216	519	293	145	545
CAMROSE	381	274	0	182	97	521	553	463	453	245	429	129	348
DRUMHELLER	263	138	182	0	279	703	735	547	282	416	247	165	530
EDMONTON	401	294	97	279	0	439	456	366	509	251	526	148	250
FT. MCMURRAY	840	733	521	703	439	0	752	796	948	587	931	587	436
GR. PRAIRIE	682	720	553	735	456	752	0	397	935	701	982	586	318
JASPER	287	412	463	547	366	796	397	0	626	613	703	413	464
LETHBRIDGE	342	216	453	282	509	948	935	626	0	605	168	360	760
LLOYDMINSTER	626	519	245	416	251	587	701	613	605	0	480	374	496
MEDICINE HAT	419	293	429	247	526	931	982	703	168	480	0	409	777
RED DEER	253	145	129	165	148	587	586	413	360	374	409	0	399
SLAVE LAKE	652	545	348	530	250	436	318	464	760	496	777	399	0

Table 5.1. Distances between cities of Alberta, measured in kilometres.

The target ratings that the network was trained on were derived by taking the shortest distance in kilometers between two cities and converting this distance to a rating. These ratings were based on the following scheme: If the distance to be determined was between a city and itself, it was assigned a rating of 0. If the distance between the two cities was less than 100 kilometres, then it was assigned a value of 1; if the distance was between 100 and 199 kilometres, then it was assigned a value of 2; if the distance was between 200 and 299 kilometres, then it was assigned a value of 3; and so on up to a maximum value of 10 which was assigned to distances of 900 kilometres or more. The complete set of ratings that were used is provided in Table 5.2.

	BANFF	CALGARY	CAMROSE	DRUMHELLER	EDMONTON	FT. MCMURRAY	GR. PRAIRIE	JASPER	LETHBRIDGE	LLOYDMINSTER	MEDICINE HAT	RED DEER	SLAVE LAKE
BANFF	0	2	4	3	5	9	7	3	4	7	5	3	7
CALGARY	2	0	3	2	3	8	8	5	3	6	3	2	6
CAMROSE	4	3	0	2	1	6	6	5	5	3	5	2	4
DRUMHELLER	3	2	2	0	3	8	8	6	3	5	3	2	6
EDMONTON	5	3	1	3	0	5	5	4	6	3	6	2	3
FT. MCMURRAY	9	8	6	8	5	0	8	8	10	6	10	6	5
GR. PRAIRIE	7	8	6	8	5	8	0	4	10	8	10	6	4
JASPER	3	5	5	6	4	8	4	0	7	7	8	5	5
LETHBRIDGE	4	3	5	3	6	10	10	7	0	7	2	4	8
LLOYDMINSTER	7	6	3	5	3	6	8	7	7	0	5	4	5
MEDICINE HAT	5	3	5	3	6	10	10	8	2	5	0	5	8
RED DEER	3	2	2	2	2	6	6	5	4	4	5	0	4
SLAVE LAKE	7	6	4	6	3	5	4	5	8	5	8	4	0

Table 5.2. The distances from Table 5.1 converted into ratings on a 0 to 10 scale.

This set of judgments was considered metric in nature for four reasons: 1) the ratings matrix was based on physical distances between geographic locations, 2) the minimality principle was enforced by requiring that the distance between a place and itself be rated as 0, 3) the symmetry principle was enforced by rating City 1 to City 2 the same distance as City 2 to City 1, and 4) the triangle inequality principle was adhered to by ensuring that for any given triad of cities the distance between any two cities was shorter than the distance between all three. As confirmation, a subroutine programmed in Excel using Visual Basic was used to check that all possible triad combinations of all thirteen cities adhered to the triangle inequality principle. Finally, as additional corroboration, multidimensional scaling was applied to the ratings matrix. The effectiveness of MDS rests on the assumption that the data being analyzed has properties that adhere to the restrictions of metric space. Hence, an MDS solution that accounts for a large amount of the variance in a data set would suggest that data has metric properties. For the ratings data, a two-dimensional MDS solution accounted for 99.4% of the variance. When plotted on a scatterplot, the points from this solution created a configuration very like the map of Alberta (see Figure 5.1).

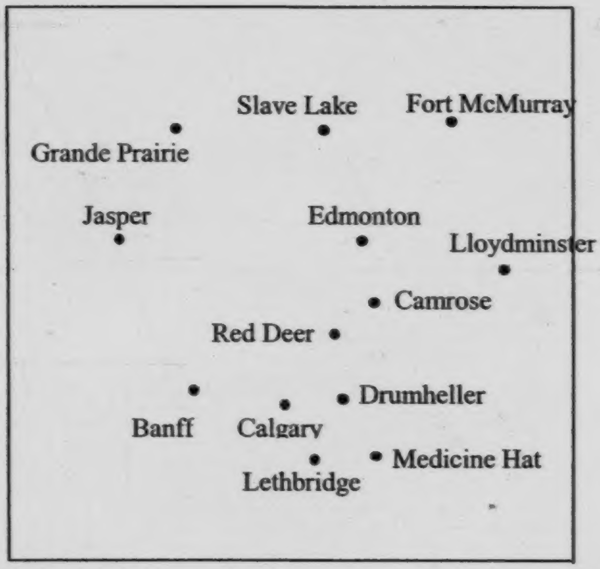

Figure 5.1. The locations of 13 Albertan cities taken from the MDS analysis of the ratings matrix of Table 7. The positions of the cities are very similar to their locations on a map of Alberta.

Network Architecture

The network that was trained to make the distance ratings was a network of value units that had 10 output units, 6 hidden units, and 13 input units as illustrated in Figure 5.2.

Input unit representation. Each input unit represented one of the thirteen place names. Pairs of places were presented as stimuli by turning two of the input units on (that is, by activating them with a value of 1). For example, to ask the network to rate the distance between Banff and Calgary, the first input unit would be turned on (representing Banff), as would the second input unit (representing Calgary). All of the other input units would be turned off (that is, were activated with a value of 0). This unary representational scheme was chosen because it contains absolutely no information about the location of the different places on a map of Alberta. In other words, the input units themselves did not provide any metric information that the network could use to perform the ratings task.

Hidden units. The layer of hidden units in this network consisted of six value units. Preliminary testing revealed this was the least number of hidden units that could be used and still allow the network to learn the correct output responses, that is, learn the ratings task accurately. Previous research has suggested that using the least possible hidden units results in a network that is easier to interpret (Berkeley et al., 1995).

Output unit representation. Ten output units were used to represent the network's rating of the distance between the two place names presented as input. The output units were also value units. To represent a rating of 0, the network was trained to turn all of its output units off. To represent any other rating, the network was trained to turn on one, and only one, of its output units. Each of these output units represented one of the ratings from 1 to 10. For example, if the network turned output unit 5 on, this indicated that it was making a distance rating of 5.

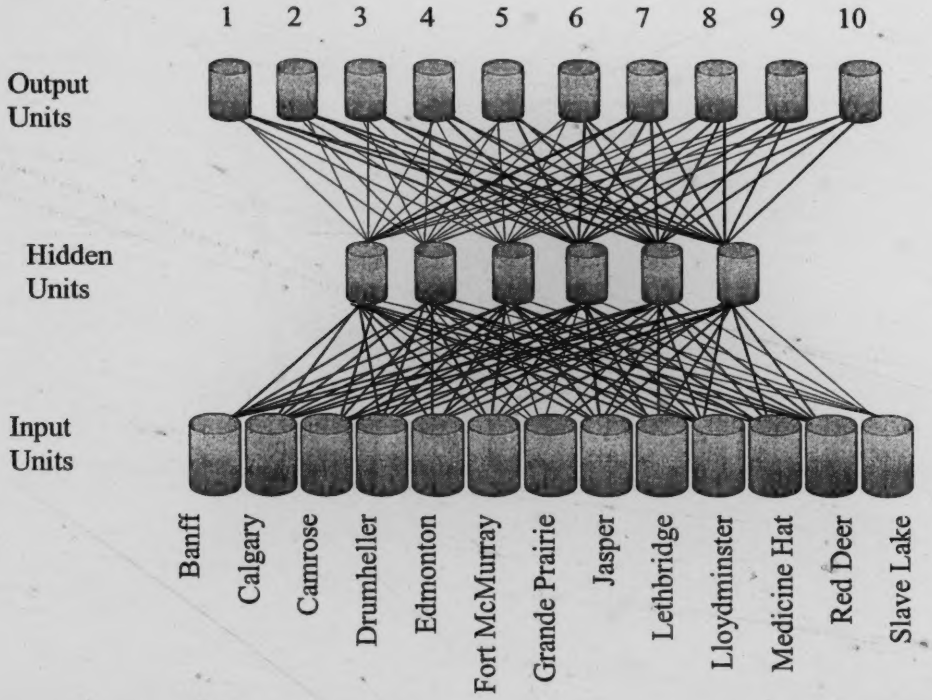
1
2
3
4
5
6
7
8
9
10
Output Units
Hidden Units
Input Units
Banff
Calgary
Camrose
Drumheller
Edmonton
Fort McMurray
Grande Prairie
Jasper
Lethbridge
Lloydminster
Medicine Hat
Red Deer
Slave Lake

Figure 5.2. A multilayer perceptron. The environment activates input units, which in turn produce activity in a layer of hidden units. The hidden units then produce a response by sending signals that produce activity in the output units. Links between layers are weighted connections; the values of the weights are determined by having the network learn through examples. Given two input city names, the network was trained to rate the distance between the cities on a map of Alberta.

Network Training

The network was trained using a variation of the generalized delta rule that has been developed for networks of value units (Dawson & Schopflocher, 1992). Prior to training, all of the connection weights were randomly assigned values ranging from –0.10 to +0.10. The biases of processing units (i.e., the means of the Gaussian activation functions, which are analogous to thresholds) were randomly assigned values ranging from –0.50 to +0.50. The network was trained with a learning rate of 0.10 and zero momentum. During each sweep of training, each of the 169 stimuli was presented to the network. The learning rule was used to update connection weights in the network after each stimulus presentation. Prior to each sweep of training, the order of stimulus presentation was randomized.

After 10,907 sweeps, the network was able to generate a correct rating for each of the 169 patterns. Correct responses for the output units were operationalized as an activation value of 0.90 when the desired activation was 1.00 and as an activation value of 0.10 or lower when the desired activation was 0.00.

5.6.2 Results

The network learned the task of producing the correct ratings from 0 -10 between all pairs of the thirteen cities in Alberta.

Given that the output ratings produced by the network adhere to metric properties, the question of interest focuses on the type of internal representations that the network relied on to generate this metric behaviour. What information is used by the hidden units to represent the spatial configuration of the map of Alberta?

Relating The Map Of Alberta To Hidden Unit Connection Weights

As described earlier, a PDP network learns by making adjustments to the connection weights between units in the network. These adjustments are based on the errors between initial responses and target responses. Therefore, the possibility that the network's representations of distance are contained in the connection weights is likely.

Neuroscience evidence regarding the function of place cells in the hippocampus establishes the premise that individual processing units may be responsible for representing specific spatial regions of a given environment. Therefore, the first analysis of the internal representations of the network is based on the supposition that each hidden unit may be calculating distance by establishing its own position in the map of Alberta. Each hidden unit then, based on its connection weights, calculates a distance to each of the cities relative to it's own position. If this hypothesis is correct, an optimal position on the map of Alberta for each hidden unit should be identifiable where there is a substantial correlation between the unit's connection weights and the distances from each city to the hidden unit location.

The Solver tool in Microsoft Excel provided an objective method for finding this location for each hidden unit. The Solver tool is a program that manipulates the values of specified cells in a spreadsheet as it attempts to optimize the value of one other cell (Orvis, 1996). In this case, the Solver tool was used to optimize the correlations between the map distances and connection weights for each hidden unit.

An Excel spreadsheet was created for each hidden unit that contained a column for the latitude and the longitude for each of the thirteen cities and a column for the final connection weight for each city for that unit. To begin the analysis, the hidden unit was first assigned an initial latitude and longitude that positioned it in the center of the map of Alberta. The spreadsheet was designed to calculate the distance for each of the cities to the hidden unit in this starting position. These 13 distances were then correlated with the 13 connection weights feeding into the hidden unit. This correlation was the value that the Solver tool optimized. The Solver tool searched for the latitude and longitude of the hidden unit that produced the most extreme correlation between distances and connection weights. For each hidden unit, a map location for the unit was found that produced strong correlations between city distances to that location and the hidden unit's incoming connection weights (see Table 5.3).

These correlations suggest that each of the connection weights leading into a hidden unit represent the distance from that unit to a particular city when the hidden unit holds it's own specific location on the map. Figure 5.3 illustrates the map of Alberta with both the actual locations of the 13 cities and the calculated locations of the 6 hidden units.

Hidden Unit	Hidden Unit Latitude	Hidden Unit Longitude	Correlation Between Map Distances And Incoming Weights
H0	51.72	113.55	0.88
H1	50.84	113.63	0.59
H2	50.88	117.70	0.72
H3	53.39	115.05	-0.54
H4	53.91	113.42	0.79
H5	51.17	115.57	-0.48

Table 5.3. Results of relating Alberta map distances between cities and hidden units to the values of the connection weights feeding into the hidden units in Simulation 1.

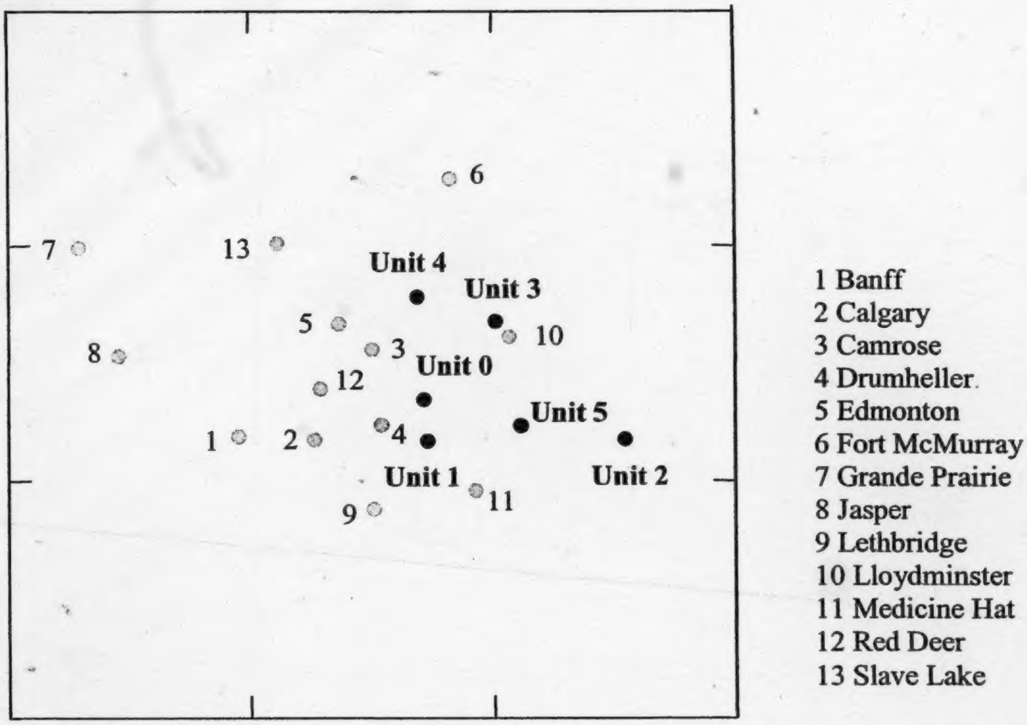

Figure 5.3. The locations of the six hidden units from the first simulation on the map of Alberta. These locations optimized the correlations between hidden unit connection weights and distances from the cities to the hidden units.

Relating Connection Weights To Hidden Unit MDS Spaces

Although there were clear correlations between the hidden unit connection weights and the actual map distances, these correlations were not as strong as anticipated. There are two explanations to consider as to why the hidden units would have a distorted representation of the map. First, conceptually, if connection weights leading into a hidden unit represent distance, then these distances are dramatically transformed by the Gaussian activation function of the hidden unit when connection weight signals are converted into hidden unit activity. This kind of conversion could be expected to result in a distortion of the actual map configuration.

Second, for any given pair of cities, the activation value for a particular hidden unit could be thought of as that unit's distance rating. However, if we examine hidden unit activity to various pairs of cities, we can see that a hidden unit's activation value does not correspond uniquely to a particular distance. For example, for hidden unit 2, when the network is asked to rate the distance between Red Deer and Jasper, this unit generates an activation value of 0.69. The actual distance between Jasper and Red Deer on the Alberta map is 413 km. However, for Edmonton and Lloydminster, which are much closer together on the map (251 km), hidden unit 2 generates an almost identical activation value of 0.71. If the network is generating the correct ratings for all pairs of cities, then what is the basis for this discrepancy?

The previous analyses revolve around the assumption that the hidden units are spatially organized and that their organization is based on the map of Alberta. This assumption ignores the possibility that the hidden units may be spatially organized around some other space. If we consider that the hidden units are indeed spatially organized but are organized around a different space than the actual map, the question is what could be the origin of that space?

The distance information contained within the entire network is represented in the connection weights between units. However, the distance information contained in a single hidden unit can only be represented by that unit's activation value. It is reasonable then to consider that the activation value of a particular hidden unit represents that hidden unit's rating of a particular pair of cities. For each hidden unit, every possible pairing of the thirteen cities would result in a matrix of 169 activation values. If these values are considered equal to ratings, MDS can be applied to this data matrix to discover the structure of the space that underlies each hidden unit's behaviour. It would then be possible to investigate how the properties of that space and the connection weights of the hidden unit may be associated.

Therefore, in this second analysis, a 13 X 13 "activity matrix" for each hidden unit was created, in which each row and each column corresponded to an Albertan city, and each cell was the hidden unit's activation value for a particular pair of cities. MDS was used to analyze each of these activity matrices, in order to create a space that represented each hidden unit's perspective on the map. For each hidden unit, a two-dimensional MDS solution accounted for almost all of the variance in the data. This solution generated an R^2 of 0.98 for hidden unit 0, 1.00 for hidden unit 1, 0.99 for hidden unit 2, 1.00 for hidden unit 3, 0.99 for hidden unit 4, and 0.97 hidden unit 5. However, each of the two-dimensional plots based on these solutions were different and none resembled the map of Alberta at all. Nonetheless, what this MDS analysis did accomplish was to provide a set of new city coordinates that was customized for each hidden unit.

How do these MDS city coordinates, relate to the hidden unit connection weights? Excel Solver was again used to optimize the correlations between each unit's connection weights and the distances from each city to the unit but in these analyses the new city coordinates based on each MDS space were used instead of the actual city coordinates. In this way, for each hidden unit a location was found in the MDS space that produced a near perfect correlation between distances and connection weights, as is reported in Table 5.4. This location represents the point of reference for all the other points in the space. This type of spatial referencing is referred to as allocentric, meaning, " centred on another" (Grush, 2001).

Hidden Unit	X-Coordinate Of Hidden Unit	Y-Coordinate Of Hidden Unit	Correlation Between MDS Distances And Connection Weights
H0	1.53	-1.88	-0.95
H1	-0.07	0.10	0.96
H2	-1.33	0.31	-0.88
H3	-0.13	-0.14	0.93
H4	0.07	0.16	-0.95
H5	3.20	0.19	-0.96

Table 5.4. Results of relating connection weights to city distances from the MDS solutions obtained from the activity matrix for each hidden unit in Simulation 1. The table provides the maximum correlation, as well as the coordinates of the hidden unit in the space that produces this maximum correlation.

Coarse Coding From Hidden Unit Activations To Distance Ratings

We can see from these analyses that the network was able to form an internal representation that encoded the spatial relationships between the thirteen cities on the map. This was accomplished by using the connection weights as representations of the distances between the hidden units and the cities in a two-dimensional space.

However, although there are clear relationships between the connection weights in the network and the distances on the map of Alberta, the individual hidden units themselves do not provide accurate assessments of distances between cities. In fact, at times, a similar activity value generated by a hidden unit for different pairs of cities will not produce the same output response (rating) by the network. To verify this claim quantitatively, the activity of each hidden unit was correlated with the desired rating for the input patterns. For units H0 through H5, these correlations were –0.32, 0.04, 0.04, -0.10, 0,04, and 0.16. As well, none of the two-dimensional MDS plots generated from the activity matrix for each hidden unit resembled the configuration of city locations in Alberta. Yet, ultimately, the network was able to generate the correct output ratings for each pairing of all the thirteen cities.

This is possible because of a type of representation called coarse coding. (e.g., (Churchland & Sejnowski, 1992, pp. 178-179). Essentially, coarse coding means that, unlike a processor tuned to detect a very precise feature, the individual units in the network are sensitive to a wider range of values of a particular feature. These value ranges overlap from unit to unit resulting in a final rating that is the pooled output of all the hidden units, a very accurate representation of a particular feature, in this case a distance. When the network is asked to generate a rating for a specific pair of cities, each hidden unit produces a rough estimate of the combined distance from the two cities to the hidden unit. When the estimates from all six units are pooled together, the network can produce an accurate representation of that distance. In order to solve the problem, the hidden units in the network adopted this type of coding.

If the hidden units in the network are indeed using coarse coding, then the combined activation values for all six units for each of the 169 stimuli should account for a large amount of the variance in the distance ratings. A regression analysis indicated that this was the case. Multiple linear regression was used to predict the distance ratings from the activation values produced by the six hidden units for each of the 169 stimuli presented during training. The regression equation produced an R^2 of 0.691, $F(6,163) = 60.714$, $p < 0.0001$).

Once the network was trained to solve the task, the network could account for 100% of the distance ratings due to the Gaussian activation functions of the output units, which resulted in a nonlinear transformation of the signal sent by the hidden units.

5.6.3 Discussion

In Simulation 1, a PDP network was trained to correctly rate the distances between thirteen cities in Alberta. Several analyses of the internal representations developed by the hidden units in the network indicated that these representations were metric in nature. First, a two-dimensional MDS solution accounted for most of the variance in the activity matrix for each hidden unit. Second, a virtual position for each hidden unit could be determined that produced strong correlations between the incoming connection weights for a given pair of cities and the actual map distance between that pair of cities. Third, these correlations were even higher when a customized space for each hidden unit, as determined by MDS analyses of its activation matrix, was used to examine the relationship between the connection weights and the distances within that space.

In Simulation 1, the PDP network generated accurate distance ratings because it adopted a coarse coding scheme based on an allocentric spatial reference, that is, a reference system organized around the location of the hidden unit itself. We can refer to this as *allocentric coarse coding*. To summarize this process: Each hidden unit in the network estimated the positions of cities relative to it's own virtual position in a two-dimensional space. For each pair of cities, this was accomplished by combining the distance from the hidden unit to each of the two cities being rated. While none of the six individual hidden units represented a space that resembled any portion of the map of Alberta, when the relations encoded in the six separate spaces were considered together, the network was able retrieve accurate spatial information about the map of Alberta.

5.7 Simulation 2: Nonmetric Spatial Judgments

In Simulation 1, the network was trained to make distance judgments based on representations that adhered to the properties of metric space. The resulting internal representations of the network also exhibited spatial properties. These properties were based on an unexpected referencing scheme, allocentric referencing, where each hidden unit occupied an imaginary location in the map of Alberta and calculated a total distance from itself to the two cities being rated. Coarse coding across the hidden units was responsible for the network's ability to generate accurate distance ratings even though the individual hidden units themselves produced distorted distance information.

The overall theme of this chapter is the relationship between spatial representation and behaviour, specifically, can there be disparities between the properties of a behaviour and the underlying representation that produced that behaviour? The representations of distance used in Simulation 1 followed metric principles, consequently it is not wholly unexpected that these representations resulted in metric judgments. But what happens if the information that a network is trained on violates metric properties? Is it possible for a metric representation to bring about nonmetric judgments? The coarse coding representational scheme revealed in the first simulation allows for the possibility that the metric representation of the network might be flexible enough to generate distance judgments that are not entirely metric. To test this hypothesis, a second simulation was designed to violate the minimality principle of metric space. Two questions of interest were: Will such a network learn to discern the correct distances between the thirteen places? Will the internal representation of this network also rely on coarse coding or will it require a different form of representation to solve the task?

5.7.1 Methods

In the first simulation the minimality principle was upheld by having the network learn to rate a distance as 0 whenever it was asked to rate the distance between a city and itself. In this second simulation, the minimality principle was violated. For some of the cities, the network was trained to make a judgment of 2 when rating the distance of a city to itself (Camrose, Grande Prairie, Lloydminster, Red Deer, Slave Lake). For some of the other cities, the network was trained to make a judgment of 1 when rating the distance of a city to itself (Drumheller, Fort McMurray, Lethbridge, Medicine Hat). For the remaining cities, the network was trained to make a judgment of 0 when rating the distance of a city to itself (Banff, Calgary, Edmonton, Jasper). This was accomplished by inserting the values listed above into the diagonal entries of Table 1b. These values were chosen to reflect what is known about human similarity judgments. Subjects are more likely to give higher self-similarity ratings (which would be converted into shorter distance ratings) to more familiar items (Tversky, 1977). Based on such findings, a rating of 0 for the four most familiar Alberta cities was used, a rating of 2 for the five smallest and least familiar cities, and a rating of 1 for the remaining four cities.

Network Architecture and Training

The input unit and output representations in the second simulation were identical to those used in the first simulation. However, a seventh hidden unit was required for the network to converge on a solution.

As before, the network was trained using the modified generalized delta rule developed by Dawson and Schopflocher (1992). Connection weights and biases were randomly started in the range from –0.50 to +0.50. The network was trained with a learning rate of 0.01 and with zero momentum; the order of pattern presentation was randomized every sweep through the training set. After 2,057 sweeps the network converged on a solution for the problem.

5.7.2 Results

The network used in Simulation 2 was able to solve the spatial configuration task even though the training information did not entirely adhere to metric principles. Beyond the network's ability to complete the task, an interesting issue is the nature of the internal representations that enabled it to do so. Do the relationships found in Simulation 1 between the connection weights, the activation values and the map distances hold true for Simulation 2? Or did the second network develop an entirely new way of representing the map distances? Given that the second network required an additional hidden unit to process the information, there are two possibilities regarding its internal representations. First, it is possible the first six units in the network processed the metric information just like the six units in Simulation 1, while the seventh unit behaved in a more specialized manner, dealing with the non-metric information on its own. Second, the network in Simulation 2 could have processed the information across all seven units using the allocentric coarse coding scheme outlined for Simulation 1. To investigate which explanation of internal representation accounted for the results of Simulation 2, the same analyses were conducted as those for Simulation 1.

Relating the Map of Alberta to Hidden Unit Connection Weights

As with Simulation 1, the Excel Solver tool was used to identify the location for each hidden unit on the map of Alberta that optimized the correlation between each hidden unit's connection weights and the distances between the unit and each of the thirteen cities. Again, a position was found for each hidden unit that produced a considerable correlation between the hidden unit connection weights and the map distances (Table 5.3). A quick comparison of the correlations generated for the six hidden units in Simulation 1 (Table 5.3) and the correlations generated for the first six hidden units in Simulation 2 (Table 5.5), suggests that no hidden unit in Simulation 2 is behaving qualitatively different from any other hidden unit in that network. This indicates that none of the hidden units in Simulation 2 are acting as a *specialist* to deal with the nonmetric information introduced to this network but rather, as in Simulation 1, the spatial information is distributed across the hidden units.

When the hidden unit positions from this analysis are illustrated on a scatterplot, none of the units share a position with a real city location. This is the case for Simulation 1 as well although the hidden unit positions for Simulation 1 (Figure 5.3) are different from those in Simulation 2 (Figure 5.4). Because the units occupy distinct positions from the cities, the distances from the units to the cities, that is, the connection weights, must always be greater than zero. This particular property of the network may be responsible for the network's ability to deal with the violation of the minimality principle because even in the metric network, a hidden unit never has a distance of zero between itself and a city.

Hidden Unit	Hidden Unit Latitude	Hidden Unit Longitude	Correlation Between Map Distances And Incoming Weights
H0	50.55	112.99	0.69
H1	54.84	115.17	0.54
H2	54.06	113.14	-0.69
H3	50.02	110.72	0.46
H4	54.51	112.20	-0.76
H5	51.17	115.57	-0.57
H6	52.75	115.44	-0.41

Table 5.5. Results of relating Alberta map distances between cities and hidden units the values of the connection weights feeding into the hidden units in Simulation 2.

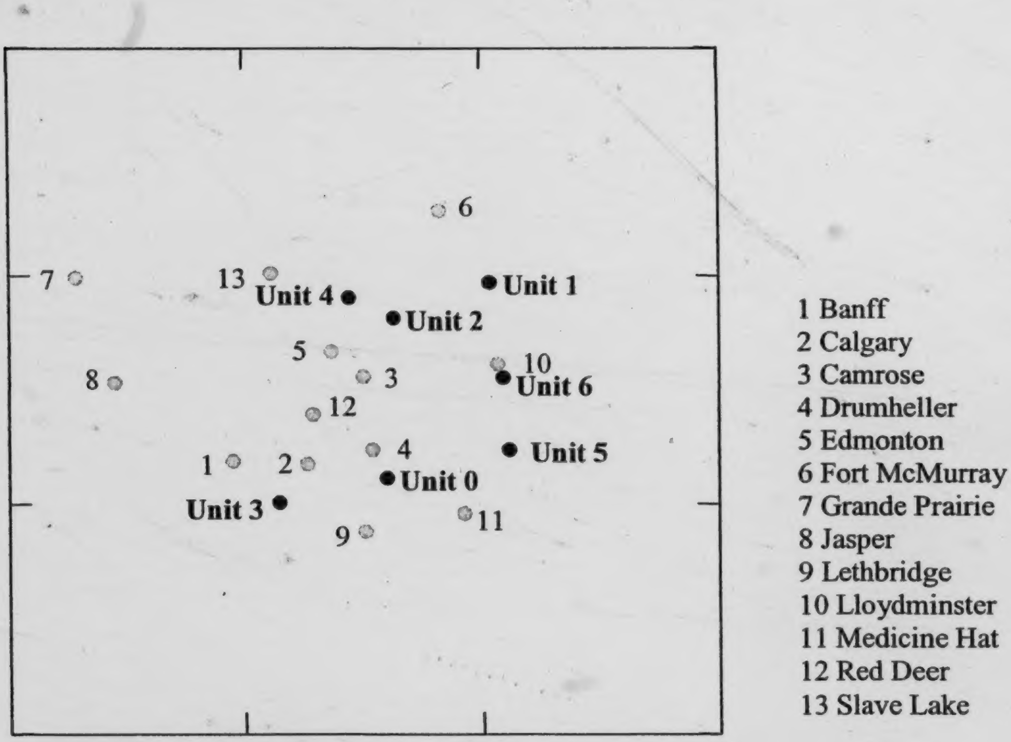

Figure 5.4. The locations of the seven hidden units from the second simulation on the map of Alberta. These locations optimized the correlations between hidden unit connection weights and distances from the cities to the hidden units.

Relating Connection Weights to Hidden Unit MDS Spaces

In accordance with Simulation 1, the correlations between the hidden unit connection weights and the map distances were sizeable but not compelling. Consequently, the second analysis used for Simulation 1 was repeated for Simulation 2. MDS was applied to the activity matrix for each of the seven hidden units to identify the coordinate of the hidden unit in its customized two-dimensional space and the city coordinates that also corresponded to that space. Again, the two-dimensional solutions accounted for close to all the variance in the data ($R^2 = 0.981$, 0.991, 0.992, 0.998, 0.992, 0.996, 0.998 for the analysis of hidden units 1 through 7). Using Solver, the location for each hidden unit was found that optimized the correlations between the connection weights and the MDS city distances. As shown in Table 11, similar to Simulation 1, these correlations turned out to be very strong.

Hidden Unit	X-Coordinate Of Hidden Unit	Y-Coordinate Of Hidden Unit	Correlation Between MDS Distances And Connection Weights
H0	-4.08	-0.98	-1.00
H1	-0.27	-0.03	0.84
H2	0.14	-0.11	0.87
H3	0.00	0.12	0.80
H4	0.36	-0.02	0.58
H5	0.38	-0.48	-0.93
H6	-0.90	-0.35	0.88

Table 5.6.
Results of relating connection weights to city distances from the MDS solutions obtained from the activity matrix for each hidden unit in Simulation 2. The table provides the maximum correlation, as well as the coordinates of the hidden unit in the space that produces this maximum correlation.

Coarse Coding from Hidden Unit Activations to Distance Ratings

For Simulation 1, coarse coding was introduced as an explanation for the network's ability to complete the distance ratings task because individual hidden units themselves did not seem responsible for generating any specific map distances. This was demonstrated by the low correlations between individual hidden unit activities and the distance ratings. However, the combined activities of all units were strongly related to the distance ratings. This was the case for Simulation 2 as well. When individual hidden unit activities were correlated with distance ratings, the correlations were weak (for units H0 through H6 respectively, –0.51, 0.22, -0.06, -0.03, -0.08, 0.06, and 0.17). However, when the activities of all seven hidden units were combined, a multiple regression analysis revealed the hidden unit activities predicted approximately 70% of the variance in the distance ratings. The regression equation produced an R^2 of 0.74, $\underline{F}\ (6,163) = 66.59$, $\underline{p} < 0.0001$). When this linear combination is transformed by the nonlinear activation functions in the output units, the result is near-perfect prediction (i.e., a converged network).

Coarse Coding for Violations of Minimality

To this point, the analyses for both Simulation 1 and Simulation 2 have produced similar results, indicating the internal representations of both networks are comparable. In addition, the second network did not seem to recruit any particular hidden unit to deal with the minimality violations introduced into a subset of the ratings. To investigate this possibility further, the modified diagonal entries were correlated with the incoming connection weights to each hidden unit. Given that these connection weights represent distance information, if one hidden unit is processing the minimality violations independently, then the correlations between the modified diagonal ratings and that unit's connection weights should be higher than those between any other weights and units. In fact, this was not the case. No individual unit produced correlations large enough to suggest it was processing the minimality violations on its own and, generally, the correlations for all seven units were comparable. The correlations that were computed for each of the seven hidden units were 0.36, -0.41, 0.22, -0.32, 0.30, -0.12, and –0.32. Thus, it appears the ratings that violate metric principles are processed in the same manner as those that adhere to metric principles, that is, across all seven units by means of coarse coding.

5.7.3 Discussion

In Simulation 2, a PDP network was trained to rate the distances between the same thirteen cities in Alberta as in Simulation 1. Unlike Simulation 1, some of the ratings the network was trained on violated a principle of metric space, that is, the minimality principle. Despite this being a more difficult task, the network was able to converge on a solution but required an additional hidden unit to do so. The question of interest for Simulation 2 was - if spatial judgments are made based on nonmetric information, will the resultant internal representations of the network be qualitatively different from those judgments based on metric information?.

Based on the same analyses as those used in Simulation 1, the second network's internal representations did not appear to be qualitatively different from those in the first network. Again, for each hidden unit, a two-dimensional MDS solution accounted for almost all the variance in its activity matrix. Strong correlations between the incoming connection weights and the actual map distances were found when an optimal virtual location was established for each hidden unit. Using a customized space derived from the activity matrix of each hidden unit, even stronger correlations were produced between the connection weights and the distances within that space.

Furthermore, indicated by comparable correlations for each hidden unit, no single hidden unit seemed to be processing the metric violations as a specialist. This suggests that the coarse coding representational scheme provides the network with the flexibility to process nonmetric information and still generate metric behaviour (e.g., ratings that follow metric principles).

5.8 Summary and Implications

The introduction of this chapter reviews three proposals of spatial mental representation involving the notion that mental representations maintain the metric properties of real-space: similarity spaces, mental imagery and cognitive mapping. These proposals are all based on the evidence that observed behaviours have metric qualities themselves. However, behavioural evidence alone has not resolved the debate. The synthetic approach provides evidence that suggests it is both. Metric representations can mediate metric behaviours but nonmetric representations can also mediate metric behaviours. The two simulations described in Chapter five not only provide evidence that both are possible but also revealed one mechanism by which such representations can be accomplished – allocentric coarse coding.

Two neural networks were trained to rate the distances between thirteen cities in Alberta. The first network was trained on ratings that adhere to the metric principles of space. Once the network was trained, examinations of its internal representations revealed a representational scheme, allocentric coarse coding, that accounted for the network's ability to learn the task even though it was never given metric input directly. Allocentric coarse coding has two essential properties. First, this representational scheme is termed allocentric because each processing unit in the network appears to represent the relations between cities as distances from itself in a virtual location to every actual city location. Second, this representational scheme is referred to as coarse coding because the final correct distance ratings are a result of the combined activities across all processing units in the network. Each hidden unit alone does not produce an accurate rating but their combined outputs, in conjunction with the Gaussian activation function, produce a correct response for every pairing of the thirteen cities. This is possible because the individual units in the network are sensitive to a wide rather than narrow range of distance values. These value ranges overlap from unit to unit resulting in a final rating that is the pooled output of all the hidden units.

The second network differed in two ways from the first: 1) the ratings the network was trained on contained values that violated the minimality principle of space, 2) the network required one additional hidden unit to process these violations. This second network was able to process these violations without any qualitative changes in its internal representations of distance from those used by the first network. No individual processing unit was responsible for processing these violations; the network appeared to use the same allocentric coarse coding scheme as that described for the first network. It is quite possible the flexibility provided by the wider range of sensitivity of each unit to the ratings values and the resultant overlapping effect between units was responsible for the network's ability to process the discrepant values.

The synthetic approach yielded a result that suggested a representation, allocentric coarse coding, that would not be apparent from behavioural evidence alone. This coding scheme has properties that explain how different representations can produce similar behaviours. The PDP simulations also illustrated that in order to process discrepancies between the properties of representation and the properties of behaviour, additional processing resources are necessary (e.g., an additional hidden unit).

These findings also have implications for the functioning of place cells in the hippocampus. The allocentric coarse coding representation of spatial information introduced in this study provides one feasible explanation for two concerns raised in the literature on the hippocampus. First, even though place cells appear to play a significant role in the instantiation of the cognitive map, there does not appear to be any regular topographic organization of place cells relative to either their positions within the hippocampus or to the positions of their receptive fields with respect to the environment (Burgess, Reece & O'Keefe, 1995; McNaughton et al., 1996). Second, place cell receptive fields are at best locally metric (Touretzky, Wan & Redish, 1994). This is because information about bearing cannot be garnered from place cell representations, and the distance between points that are more than about a dozen body lengths apart cannot be measured because of a lack of place cell receptive field overlap.

The PDP networks display the same properties that suggest place cells may not be capable of rendering totally metric representations of space. First, because the hidden units are all connected to all of the input units, the network has no definite topographic organization. Second, each hidden unit appears to be at best locally metric. While the input connections can be correlated with distances on the map, the responses of individual hidden units do not provide an accurate spatial account of the map. Yet these locally metric processors are capable of producing a configuration of the thirteen cities that accurately reflects the map of Alberta. This is possible because the network does not base its output on the behaviour of a single hidden unit but rather, via coarse coding, it generates its response from the activity of all hidden units considered simultaneously.

CHAPTER 6

6.1 CONCLUSIONS

The evidence presented in this thesis illustrates that different hypertext navigation behaviours are affected by the information provided in navigation tools. This is true for both navigation performance measures such as speed and recall, and also for the course of navigation as evidenced by page-transition patterns.

In Chapter two, a custom program for data collection was described. This program addresses several obstacles in the collection of relevant data, specifically, potential difficulties with relying on inherent mechanisms within traditional webrowsers. For example, browser history lists make use of a variety of procedures for recording page access. To optimize the history list, some methods erase prior revisits and only retain the most recent instance of access (e.g., for further description of history lists see Tauscher & Greenberg, 1997). By creating a custom application, a local log file of users' movements can be constructed to record every single move with no elimination of repeated page accesses that can happen with the typical history lists of standard browsers. By having tracking procedures local to each individual workstation, that is, creating local logs instead of server logs, we can avoid the above-mentioned problems.

Regarding methodology, the distance-like ratings task is a unique task developed to tap into users mental representations of the document space. The task is based on the premise that a psychological distance can be a measure of semantic relatedness. By asking users to provide distances between numerous pairs of pages, we can gain a sense of what information was important for organizing material in memory.

The novel methods introduced in this thesis have promise for investigating additional, related areas of hypertext use. For instance, different user characteristics may alter the influence of certain types of information. As a specific example, the navigation tools, data collection program and ratings task may be used to study the effects of reading skill on users' interactions with given interface elements.

As reported in Chapter three, for performance measures, including both types of information in a single navigation tool did not enhance user speed or economy over tools with either spatial or hierarchical information alone. However, all three of the navigation tools with spatial and/or hierarchical information produced more efficient search than the alphabetical list. The recall measure was the only performance measure that indicated spatial information combined with hierarchical information enhances performance over hierarchical information alone.

Evidence for the differential impact of navigation tool information on users' path patterns was provided in Chapter four. The examination of page-transition patterns shows that, generally, users' patterns of page-transitions seem to reflect the most prevalent organization in a given navigation tool whether that is conceptual, alphabetical or spatial, although transition patterns do not indicate an absolutely linear usage of the tool. Finally, as evidenced by users' ratings of the distances between pages, spatial information within a navigation tool seems to influence users' mental representations of the document space. This is mediated by the users' reliance on the navigation tool to make navigation decisions and their subsequent route through the content space of the document. Spatial information plays a unique role in this process.

The behavioural data in this thesis points toward a complicated interaction between spatial and conceptual information. The combination of spatial and conceptual information does not produce a summative effect on performance. Relationships do exist between behaviour and the information in the navigation tool but this is not a perfect relation. Spatial information is implicated as contributing to the development of a continuous relationship between an initial hypothetical organization for the document material, the behaviour elicited by such an organization and a final mental representation of the material. However, even though spatial attributes of the navigation tool may contribute to some kind of spatial organization of the material, that spatial representation may not be comparable to a spatial map. Rather, the representation may incorporate information such as that suggested by Pylyshyn (in press), functional descriptors that represent spatial relationships among elements. Such a possibility led to a departure from human data to an investigation of the relationship between behaviour and representation at a more abstract, theoretical level.

The nature of spatial mental representation was investigated in Chapter five using a synthetic approach which yielded evidence beyond that uncovered by behavioural analyses. An in-depth analysis of two neural networks trained to complete a spatial task revealed that different underlying representations can produce similar behaviours but additional processing resources are required to do so.

The implications of this network study for the applied issue of hypertext navigation are twofold: 1) discrepancies that appear between representation and behaviour may require more cognitive resources to produce usable mental representations of the document, resulting in longer processing times and requiring experience with more inputs (e.g., accessing more pages of the document) and, 2) a synthetic approach presents a promising opportunity to explore navigation data from a new perspective. For example, if the underlying representations of users' represent an integrated space constructed from multiple sources of varying types of information, using the MDS coordinates derived from the human data as input for network processing could be revealing. In this sense, future research in the area of hypertext use might do well to incorporate connectionist representations of navigation behaviours to facilitate our understanding of the properties and functions of underlying mental representations.

The broader implications of this network study reach into areas of general human cognition and neuroscience. Specifically, it speaks to the issue of spatial representation in the place cells of the hippocampus. Place cells have been found to involve some properties of metric space, for example, the Gaussian tuning curves seem to calculate distance from environmental elements. However, it has been argued that place cell circuitry by itself does not provide a cognitive map that can be considered metric for two reasons: 1) place cells are not organized topographically, and 2) place cell receptive fields may be only locally metric. Given these constraints, not all types of spatial information (e.g., information about bearing) could be encoded in the hippocampus neural circuitry with a strictly metric representation. The results of the network studies in this thesis suggest a representational scheme that might allow for a different type of spatial representation. In some respects the network is metric, hidden units are like place cells and distances are represented in connection weights. However, the network representation also has new unexpected properties that are not metric. Through allocentric coarse coding, processing units calculate distances based on their own optimal location in a new type of space and the processing of information occurs in a distributed fashion across all processing units. This kind of representation provides one explanation as to why topographical organization of place cells may not be necessary and also explains how metric information encoded locally by place cells may be able to contribute to a more global configuration of a given environment.

6.2 Impact to Practitioners

As stated in the introduction to this thesis, the intention of this work was to provide a theoretical investigation into psychological aspects of hypertext use. Given the focus and scope of this thesis, advocating particular design practices based on such a preliminary study would be premature. However, from an applied standpoint, two outcomes may be relevant to practitioners. The performance data in chapter three indicates that providing navigation tools that convey conceptual relationships, whether that is through explicit hierarchies or through spatial formats with conceptual information embedded within them, produces better performance that an alphabetical presentation. The data also suggests that, for most user subgoals, combining spatial and hierarchical information into one navigation tool does not produce a summative effect on performance. Hence, for navigation tool design, designers may be better off focusing on choosing one type of information and making that highly salient to the user.

6.3 Limitations

This thesis represents a preliminary investigation into the cognitive processes involved in hypertext use with a particular focus on mental representation. This requires the modification of established psychological methodologies to the particular environment of hypertext. The ratings task was devised to tap into user mental representations of the document space. The usefulness of the ratings data relies on the assumption that the ratings are a reasonable reflection of that underlying mental organization of the material. Additional studies are needed to ascertain if this is a valid assumption.

As mentioned earlier, the format of the navigation tools in this thesis represent only one option for presenting spatial information. Additional studies using various spatial contexts are required to fully comprehend the role of spatial processing in a hypertext environment.

Regarding the network studies, as with all simulations of human cognition some would argue that neural networks cannot be assumed to process information in a comparable fashion to human beings. I would address that argument with two statements: First, with traditional behavioural analyses we are also forced to make assumptions about relationships between mental process and behavioural outcomes. That is, we can never be entirely sure that our assessment of behavioural outcomes reflects actual mental processing. Second, the spatial task the network was trained on is not meant to be offered as a task equivalent to the hypertext navigation task completed by human participants but rather as an example on a theoretical level of how such outcomes of behavioural/representational discrepancy could possibly come to pass. With these two thoughts in mind, the neural networks in this thesis function as complimentary tools to behavioural paradigms that provide the additional option of exploring internal representations.

6.4 Future Research

Hypertext navigation is a complex task involving many intermediate sub goals, each of which may require different types of support. Spatial information does appear to impact some of these subgoals or intermediate processes but not all. Therefore, the spatial metaphor as it has been previously defined may be an appropriate metaphor for some of these sub goals but as an overall theoretical explanation of the cognitive processes involved in hypertext navigation it is incomplete. Further research is required to identify the interactions between sub goals, cognitive processing and diverse ambient factors in a hypertext environment. However, one point to consider that arises from the network study is that spatial representations may be very important for hypertext navigation but not the kinds of spatial representations implied by the spatial metaphor as it stands. An alternative, unusual spatial representation, not solely based on the traditional notions of space, may be responsible for such a relationship. This is an important theme for future research.

The nature of human interaction with physical spaces has been extensively studied in multitudes of spatial contexts. To gain a useful understanding of cognition in hyperspace, particularly as it applies to spatial processing, the study of mental processing as it relates to numerous spatial layouts in hypertext is also required.

References

Acredolo, L.P. (1978). Development of spatial orientation in infancy. Developmental Psychology, 14, 224-234.

Andris, J.F. (1996). The relationship of indices of student navigational patterns in a hypermedia geology lab simulation to two measures of learning style. Journal of Educational Multimedia and Hypermedia, 5 (3/4), 303-315.

Ballard, D. (1986). Cortical structures and parallel processing: Structure and function, Behavioural and Brain Sciences, 9, 67-120.

Beasley, R.E. & Waugh, M.L. (1995). Cognitive mapping architectures and hypermedia disorientation: An empirical study. Journal of Multimedia and Hypermedia, 4, 239-255.

Beasley, R.E., & Waugh, M. L. (1997). Predominant initial and review patterns of navigation in a fully constrained hypermedia hierarchy: An empirical study. Journal of Educational Multimedia and Hypermedia, 6 (2), 155-172.

Beasley, R.E. & Vila, J.A. (1992). The identification of navigation patterns in a multimedia environment: A case study. Journal of Educational Multimedia and Hypermedia, 1 , 209-222.

Berkeley, I.S.N., Dawson, M.R.W., Medler, D.A., Schopflocher, D. P. & Hornsby, L. (1995). Density plots of hidden value unit activations reveal interpretable bands, Connection Science, 7, 167-186.

Bieber, M., Vitali, F., Ashman, H., Balasubramanian, V. & Oinas-Kukkonen, H.

(1997). Fourth generation hypermedia: Some missing links for the World Wide Web. International Journal of Human-Computer Studies, 47, 31-65.

Blumenthal, L.M. (1953). Theory and applications of distance geometry. London: Oxford.

Boechler, P.M., Dawson, M.R.W. & Boechler, K.R. (in press). An introduction to custom WebBrowsers for the qualitative study of hypertext navigation. Journal of Educational Multimedia and Hypermedia.

Boechler, P.M. & Dawson, M.R.W. (in press). The effects of navigation tool information on hypertext navigation behaviour: A configural analysis of page-transition data. Journal of Educational Multimedia and Hypermedia.

Borg, I., & Groenen, P. (1997). Modern multidimensional scaling: Theory and applications. New York: Springer-Verlag.

Braitienberg, V. (1984). Vehicles: Explorations in Synthetic psychology. Cambridge, MA: MIT Press.

Burgess, N., Reece, M. & O'Keefe, J. (1995). Spatial models of the hippocampus. In M.A. Arbib (ed.), The Handbook of brain Theory and Neural Networks. Cambridge, MA: MIT Press.

Campagnoni, F.R. & Ehrlich, K. (1990). Getting there when you don't know where ' there' is: Navigational strategies in a hypertext help system. SIGCHI Bulletin, 21,3,17-18.

Canter, D., Rivers, R. & Storrs, G. (1985). Characterizing user navigation through complex data structures. Behaviour and Information Technology, 45, 2, 93-102.

References

Carroll, J.D. & Wish, M. (1974). Multidimensional perceptual models and measurement methods. In E.C. Carterette & M.P. Friedman (Eds.), Handbook of Perception. New York: Academic Press.

Carroll, J.M., Mack, R.L. & Kellogg, W.A. (1988). Interface metaphors and interface design. In M. Helander (ed.), Handbook of Human-Computer Interaction. North-Holland: Elsevier Science Publishers.

Chen, C. & Rada, R. (1996). Interacting with hypertext: A meta-analysis of experimental studies. Human-Computer Interaction, 11, 125-156.

Cheng, K., & Spetch, M.L. (1998). Mechanisms in landmark use in mammals and birds. In Healy (ed.), Spatial Representation in Animals. Oxford: Oxford University press.

Chou, C. & Lin, H. (1997). Navigation maps in a computer-networked hypertext learning system. Paper presented at the annual meeting of Association for Educational Communications & Technology, Albuquerque, New Mexico, USA.

Churchland , P.S., & Sejnowski, T.J. (1992). The Commputational Brain. Cambridge, MA: MIT Press.

Conklin, J. (1987). Hypertext: An introduction and survey. Computer, IEEE Computing Society, 20:9, 17-41 (September).

Cornell, E.H. & Heth, C.D. (1983). Spatial cognition: Gather studies used by preschool children. Journal of Experimental Child Psychology, 35, 93-110.

Cornell, E.H., Heth, C.D. & Alberts, D. (1994). Place recognition and wayfinding by children and adults. Memory and Cognition, 22, 6, 633-643.

References

Cummins, R. (1983). The nature of psychological explanation. Cambridge, MA: MIT Press.

Dawson, M.R.W. (1998). Understanding Cognitive Science, Oxford, UK: Blackwell.

Dawson, M.R.W., & Harshman, R. A. (1986). The multidimensional analysis of asymmetries in alphabetic confusion matrices: Evidence for global-to-local and local-to-global processing. Perception & Psychophysics, 40 (6), 370-383.

Dawson, M.R.W., & Medler, D.A. (1996). Of mushrooms and machine learning: identifying algorithms in a PDP network, Canadian Artificial Intelligence, 38, 14-17.

Dawson, M.R.W., Medler, D.A. & Berkeley, I.S.N. (1997). PDP networks can provide models that are not mere implementations of classical theories, Philosophical Psychology, 10, 25-40.

Dawson, M.R.W., Medler, D.A., McCaughan, D.B., Willson, L., & Carbonaro, M. (2000). Using extra output learning to insert a symbolic theory into a connectionist network, Minds and Machines, 10, 171-201.

Dawson, M.R.W., & Schopflocher, D.P. (1992). Modifying the Generalized Delta rule to train networks of nonmonotonic processors for pattern classification, Connection Science, 4, 19-31.

Dee-Lucas, D. & Larkin, J.H. (1995) Learning from electronic texts: effects of interactive overviews for information access. Cognition and Instruction, 13, 3, 431-468.

References

Dias, P. & Souza, A.P. (1997). Understanding navigation and disorientation in Hypermedia learning environments. Journal of Multimedia and Hypermedia, 6, 2, 173-185.

Dillon, A., McKnight, C. & Richardson, J. (1990). Navigation in hypertext: A critical review of the concept. Human-Computer Interaction – INTERACT'90. In D. Diaper et al. (Eds.), North-Holland: Elsevier Science Publishers.

Edwards, D.W. & Hardman, L. (1989). 'Lost in hyperspace': Cognitive mapping navigation in a hypertext environment. In R.A. McLeese (Ed.), Hypertext: Theory into Practice. London:Ablex.

Foss, C.L. (1989). Detecting user's lost: Empirical studies on the browsing of hypertext. Technical Report, No 972, INRIA, Sophia- Antipolis.

Gallistel, C.R. (1990). The organization of learning. Cambridge, MA: MIT Press.

Golledge, R.G. (1994). Primitives of spatial knowledge. In T. Nyerges (Ed.), Cognitive Aspects of Human-computer Interaction for Geographic Information Systems – NATO/ASI Series: Kluwer.

Golledge, R.G., Ruggles, A.J., Pelligrino, J.W. & Gale, N.D. (1993). Integrating route knowledge in an unfamiliar neighborhood: Along and across route experiments. Journal of Environmental Psychology, 13, 293-307.

Golledge, R.J. & Stimson, R.J. (1997). Spatial Behaviour: A geographic perspective. New York: Guilford Press.

References

Gray, S.H. (1990). Using protocol analyses and drawings to study mental model construction during hypertext navigation. International Journal of Human-Computer Interaction, 2, 128-133.

Grush, R. (2000). Self, world and space: The meanings and mechanisms of ego- and allocentric spatial representation, Brain and Mind, 1, 59-92.

Guttman, L. (1971). Measurement as structural theory. Psychometrika, 36, 465-506.

Gygi, K. (1990). Recognizing the symptoms of hypertext… and what to do about it. In B. Laurel (ed.), The Art of Human-Computer Interface Design (Addison-Wesley, Reading), 279-288.

Hammond, N. & Allinson, L. (1987). The travel metaphor as design principle and training aid for navigating around complex systems. In D. Diaper & R. Winder (Eds.), People and Computers III: Proceedings of the Third Conference of the British Computer Society, Human-Computer Interaction Specialist Group, Cambridge: Cambridge University Press.

Hart, R.A. & Moore, G.T. (1973). The development of spatial cognition: A review. In R.M. Downs & D. Stea (Eds.), Image and Environment. Chicago: Aldine.

Hirtle, S.C. & Hudson, J. (1991). Acquisition of spatial knowledge for routes. Journal of Environmental Psychology, 11, 335-345.

Holyoak, K.J. & Thagard, P. (1996). Mental Leaps: Analogy in creative thought. Massachusetts: MIT Press.

References

Horney, M. (1993). A measure of hypertext linearity. Journal of Multimedia and Hypermedia, 2 (1), 67-82.

Jones, W.P. & Dumais, S. (1986). The spatial metaphor for user interfaces: Experimental test of reference by location versus naming. ACM transactions on Office Information Systems, 4, 42- 63.

Jonassen, D. H. (1986). Hypertext principles for text and courseware design. Educational Psychologist, 214, 269-292.

Jonassen, D.H. (1989). Hypertext/hypermedia. Englewood Cliffs, NJ: Educational Technology Publications.

Kaplan, S. (1973). Cognitive maps in perception and thought. In R.M. Downs & D. Stea, (Eds.), Image and Environment. Chicago, ILL: Aldine, pp. 63-78.

Kerr, S.T. (1990). Wayfinding in an electronic database: The relative importance of navigational cues vs. mental models. Information Processing & Management, 26, 4, 511-523.

Kim, J. (1999). An empirical study of navigational aids in customer interfaces, Behaviour and Information Technology, 18, 213-224.

Kim, H. & Hirtle, S.C. (1995). Spatial metaphors and disorientation in hypertext browsing. Behaviour and Information Technology, 14, 4, 239-250.

Kitchin, R.M. (1994). Cognitive maps: What are they and why study them? Journal of Environmental Psychology, 14, 1-19.

References

Korthauer, R.D. & Koubek, R.J. (1994). An empirical evaluation of knowledge, cognitive style, and structure upon the performance of a hypertext task. International Journal of Human-Computer Interaction, 6, 4, 373-390.

Kosslyn, S.M. (1987). Seeing and imagining in the cerebral hemispheres: A computational approach. Psychological Review, 94, 148-175.

Kosslyn, S.M. (1994). Image and Brain. Cambridge, MA: MIT Press.

Kosslyn, S.M., Ball, T.M. & Reiser, B.J. (1978). Visual images preserve metric spatial information: Evidence from studies of image scanning . Journal of Experimental Psychology: Human Perception and Performance, 4, 46-60.

Kosslyn, S.M., Pascual-Leone, A., Felican, O., Camposano, S., Keenan, J.P., Thompson, W.L., Ganis, G., Sukel, K.E. & Alpert, N.M. (1999). The role of Area 17 in visual imagery: Convergent evidence from PET and RTMS, Science, 284, 167-170.

Kosslyn, S.M., Thompson, W.L., & Alpert, N.M. (1997). Neural systems shared by visual imagery and visual perception: A positron emission tomography study, Neuroimage, 6, 320-334.

Kosslyn, S.M., Thompson, W.L., Kim, I.J., & Alpert, N.M. (1995). Topographical representations of mental images in Area 17, Nature, 378, 496-498.

Klatzky, R.L., Loomis, J.M., Golledge, R.G., Cicinelli, J.G., Doherty, S., & Pellegrino, J.W. (1990). Acquisition of route and survey knowledge in the absence of vision. Journal of Motor Behaviour, 22, 19-43.

References

Krumhansl, C.L. (1978). Concerning the applicability of geometric models to similarity data: The interrelationship between similarity and spatial density, Psychological Review, 85, 445-463.

Krumhansl, C.L. (1982). Density versus feature weights as predictors of visual identifications: Comments on Appelman and Mayzner, Journal of Experimental Psychology: General, 111, 101-108.

Kruskal, J.B. (1964). Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. Psychometrika, 29, 115-129.

Kruskal, J. B., & Wish, M. (1978). MDS. Beverly Hills, CA.: Sage.

Kuipers, B. (1982). The " map in the head "metaphor. Environment and Behaviour, 14, 202-220.

Lakoff, G. (1994). What is metaphor? In J.A. Barnden and K.J. Holyoak (Eds.), Advances in Connectionist and Neural Computation Theory: Analogy, Metaphor and Reminding. New Jersey: Ablex.

Leidig, P.M. (1992). The relationship between cognitive styles and mental maps in hypertext assisted learning. Dissertation Abstracts International, 53, 5-A, 1372.

Leighton, J.P. (1999). An alternate approach to understanding formal reasoning: Thinking according to the inductive –coherence model. Edmonton: Unpublished Ph.D., University of Alberta.

Leventhal, L.M., Teasley, B.M., Instone, K., Rohlman, D.S. & Farhat, J. (1993). Sleuthing in Hyperholmes-super ™: An evaluation of using hypertext vs. a book to answer questions. Behaviour & Information Technology, 12,3, 149-164.

References

Liben, L. (1981). Spatial representation and behaviour: Multiple perspectives. In L.S. Liben, A.M. Patterson & N. Newcombe, (Eds.), Spatial Representation and Behaviour Across the Life Span. New York, NY: Academic Press

Lieblich, I. & Arbib, M. (1982). Multiple representation of space underlying behaviour and associated commentaries. The Behaviour and Brain Sciences, 5, 4, 627-660.

Marchionini, G. (1989). Information-seeking strategies of novices using a full-text electronic encyclopedia. Journal of the American Society for Information Science, 40, 54-66.

Mayes, T., Kibby, M. & Anderson, T. (1990). Signposts for conceptual orientation: Some requirements for learning from hypertext. In R. McLeese and C. Green (Eds.), Hypertext: State of the Art. Oxford: Intellect.

McDonald, S. & Stevenson, R. (1998). Spatial versus conceptual maps as learning tools in hypertext. Journal of Educational Hypermedia and Multimedia, 8, 1, 43-64.

McDonald, S. & Stevenson, R. (1999). Navigation in hyperspace: An evaluation of the effects of navigational tools and subject matter expertise on browsing and information retrieval in hypertext. Interacting with Computers, 10, 129-142.

McLeese, R.A. (1989). Navigation and browsing in hypertext. In R.A. McLeese (Ed.), Hypertext: Theory into Practice. London:Ablex.

McNamara, T.P., Halpin, J.A. & Hardy, J.K. (1992). Spatial and temporal contributions to the structure of spatial memory. Journal of Experimental Psychology: Learning, Memory and Cognition, 18, 555-564.

References

McNamara, T.P., Hardy, J.K. & Hirtle, S.C. (1989). Subjective hierarchies in spatial memory. Journal of Experimental Psychology: Learning, Memory and Cognition, 15, 211-227.

McNaughton, B., Barnes, C.A., Gerrard, J.L., Gothard, K., Jung, M.W., Knierim, J.J., Kudrimoti, H., Qin, Y., Skaggs, W.E., Suster, M. & Weaver, K.L. (1996). D Deciphering the hippocampal polyglot: The hippocampus as a path integration system, The Journal of Experimental Biology, 199, 173-185.

Mellet, E., Petiti, L., Mazoyer, B., Denis, M., & Tzourio, N. (1998). Reopening the mental imagery debate: Lessons from functional anatomy, Neuroimage, 8, 129-139.

Montello, D.R. (1993). A new framework for understanding the acquisition of spatial knowledge in large-scale environments. In R.G. Golledge & M.J. Egenhofer (Eds.), Spatial and temporal reasoning in geographic information systems.

Montello, D.R., & Pick, H.L. (1993). Integrating knowledge of vertically aligned large-scale spaces. Environment and Behaviour, 25, 457-484.

Nielsen, J. (1990). Hypertext and Hypermedia. San Diego, CA: Academic Press.

O'Keefe, J. & Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurones. Nature, 381, 425-428.

O'Keefe, J. & Nadel, L. (1978). The Hippocampus as a Cognitive Map. Oxford: Clarendon Press.

Orvis, W. J. (1996). Excel for Scientists and Engineers (Second Ed.). San Francisco, CA: Sybex.

References

Passini, R. (1984). Wayfinding in Architecture. New York: Van Nostrum Reinhold.

Presson, C.C. & Somerville, S.C. (1985). Beyond ego-centrism: A new look at the beginnings of spatial representation. In H.M. Wellman (Ed.), Children's Searching: The Development of Search Skill and Spatial Representation. Hillsdale, NJ: Erlbaum.

Pylyshyn, Z. (1980). Computation and cognition: Issues in the foundations of cognitive science. Behavioural and Brain Sciences, 3, 111-169.

Pylyshyn, Z. (1981). The imagery debate: Analogue media versus tacit knowledge. Psychological Review, 88, 1, 16-45.

Pylyshyn, Z. (1984). Computation and cognition: Toward a foundation for cognitive science. Massachusetts: MIT Press.

Pylyshyn, Z.W. (1993). Metaphorical imprecision and the "top-down" research strategy. In A. Ortony (Ed.), Metaphor and Thought. Cambridge: Cambridge University Press.

Pylyshyn, Z. (in press). Is the imagery debate over? If so, what was it about? Massachusetts: MIT Press.

Ransom, S., Wu, X. & Schmidt, H. (1997). Disorientation and cognitive overhead in hypertext systems. International Journal of Artificial Intelligence Tools, 6, 2, 227-253.

Reynolds, S.B. & Dansereau, D.F. (1990). The knowledge hypermap: An alternative to hypertext. Computers in Education , 14, 5, 409-416.

Rumelhart, D.E., Hinton, G.E., & Williams, R.J. (1986). Learning representations by back-propagating errors, Nature, 323, 533-536.

Sadalla, E.K., & Montello, D.R. (1989). Remembering changes in direction. Environment and Behaviour, 21, 346-363.

Sasse, M.A. (1992). Users' models of computer systems. In Y.Rogers et al. (Eds.). Models in the Mind. London: Academic Press.

Schneiderman, B. & Kearsley, G. (1989). Hypertext hands-on: An introduction to a new way of organizing and accessing information. NY: Addison-Wesley.

Schroeder, E.E. & Grabowski, B.L. (1995). Patterns of exploration and learning with hypermedia. Journal of Educational Computing Research, 13, 4, 313-335.

Shepard, R.N. (1962). The analysis of proximities: Multidimensional scaling with an unknown distance function. Part I. Psychometrika, 27, 219-246.

Shepard, R.N. & Metzler, J. (1971). The mental rotation of three-dimensional objects. Science, 171, 701-703.

Shum, S. (1990). Real and virtual spaces: Mapping from spatial cognition to hypertext. Hypermedia, 2, 2, 133-158.

Siegel, A.W. & White, S.H. (1975). The development of spatial representation of large-scale environments. In H.W. Reese (Ed.), Advances in Child Development and Behaviour, 10, New York: Academic.

Simon, H.A. (1996). The sciences of the artificial, Third edition. Cambridge, MA: MIT Press.

Smyth, M.M., & Kennedy, J.E. (1982). Orientation and spatial representation within multiple frames of reference. British Journal of Psychology, 73, 527-535.

Stanovich, K.E. & West, R.F. (1989). Exposure to print and orthographic processing. Reading Research Quarterly, 24, 402-433.

Stanton, N.A. & Baber, C. (1994). The myth of navigating in hypertext: How a "bandwagon" lost its course! Journal of Educational Multimedia and Hypermedia, 3, (3-4), 235-249.

Stanton, N.A., Taylor, R.G. & Tweedie, L.A. (1992). Maps as navigational aids in hypertext environments: An empirical evaluation. Journal of Educational Multimedia and Hypermedia, 1, 4, 431-444.

Stevens, A. & Coupe, P. (1978). Distortion in judging spatial relations. Cognitive Psychology, 10, 422-437.

Tauscher, L. & Greenberg, S. (1997). How people revisit web pages: Empirical findings and implications for the design of history systems. International Journal of Human-Computer Studies, 47, 97-137.

Tolman, E.C. (1948). Cognitive maps in rats and men. Psychological Review, 55, 189-208.

Tourangeau, R., & Sternberg, R.J. (1981). Aptness in metaphor, Cognitive Psychology, 13, 27-55.

Tourangeau, R., & Sternberg, R.J. (1982). Understanding and appreciating metaphors, Cognition, 11, 203-244.

Touretzky, D.S., Wan, H.S., & Redish, A.D. (1994). Neural representation of space in rats and robots. In J.M. Zurada, R.J. Marks and C.J. Robinson (eds.), Computational Intelligence: Imitating Life. New York, NY: IEEE Press.

References

Tsai, C.L. (1988). Hypertext: Technology, applications and research issues. Journal of Education Systems, 17, 1, 3-14.

Tversky, A. (1977). Features of similarity. Psychological Review, 84, 327-352.

Tversky, A. & Gati, I. (1982). Similarity, Separability and the Triangle inequality. Psychological Review, 89, 123-154.

Vicentes, K.J. & Willeges, R.C. (1988). Accommodating individual differences in searching a hierarchical file system. International Journal of Man-machine Studies, 29, 647-668.

Wenger, M.J. & Payne, D.G. (1994). Effects of a graphical browser on readers' efficiency in reading hypertext. Technical Communications, 41, 224-233.

Zimmerman, C.L. (1999). A network interpretation approach to the balance scale task. Edmonton: Unpublished Ph.D., University of Alberta.

Appendix A

Listing #1 - Administration Login Form Code

```
Option Explicit

Public LoginSucceeded As Boolean
Dim QuestionSelected As Boolean
Dim ConditionSelected As Boolean

Private Sub cmdOK_Click()
'****************************************************************
'Event: User clicks OK button
'Response:
'  If the password is correct
'     Record login success
'     Check if one of the navigation aid conditions is selected
'     Check if one of question wording options is selected
'     If both question wording and navigation condition have been selected
'        Close this form
'        Move to the test subjects login
'     Otherwise
'        Indicate required information is missing
' Otherwise
'     Indicate the password is incorrect
'     Set up the form for the user to re-enter password
'****************************************************************
If (txtPassword.Text = Psswrd) Then
        LoginSucceeded = True

    ConditionSelected = ((Option1(0).Value) Or (Option1(1).Value) Or (Option1(1).Value) Or

(Option1(2).Value) Or (Option1(3).Value))
        QuestionSelected = ((Option2(0).Value) Or (Option2(1).Value))
        If (QuestionSelected) And (ConditionSelected) Then
                Me.Hide
                LoginForm.Show
        Else

                MsgBox "Something's Missing!"
        End If

Else
        MsgBox "Invalid Password, try again!", , "Login"
        txtPassword.SetFocus
        SendKeys "{Home}+{End}"
End If
End Sub
```

Listing # 2- Subroutine# 1 Browser Form Code (Set-up)

```
Dim StartTime As String
Dim EndTime As String
Dim PageTime As String
Dim LastPage As String
Dim Answer As String
Dim Question As Byte
Dim Page As String
Dim InstructPage As String
Dim Done As String

Private Sub Form_Load()
'************************************************************
'Event: Form Loads
'Response:
'  Initialize start and end times for tracking user
'  Open navigation log files
'  Select instruction page based on admin options set in AdminLogin
'  Navigate to the instruction page
'************************************************************
Let EndTime = Timer
Let StartTime = Timer
Open (App.Path & AnswerFile) For Append As #3
Open (App.Path & LogFile) For Append As #2
Let Frame1.Visible = False
Let Submit.Visible = False
Let LastPage = App.Path & InstructFile
Select Case Condition
  Case "Hier", "Alpha"
        InstructPage = "\instructions2.htm"
  Case "Spatial", "Spatialhier"
        InstructPage = "\instructions1.htm"
   Case Else
        InstructPage = "\instructions1.htm"
End Select
WebBrowser1.Navigate App.Path & InstructPage
End Sub
```

Appendix B

Listing # 2- Subroutine# 2 Browser Form Code (Set-up)

```
Private Sub Continue_Click()
'************************************************************
'Event: Continue button clicked
'Response:
'  Set up form for questions
'  Select navigation aid page based on admin options set in AdminLogin
'  Navigate to the navigation aid page
'************************************************************
Question = 1
Frame1.Visible = True
Continue.Visible = False
Change_Question (Question)
Select Case Condition
   Case "Hier"
         Page = "\hierlist5.htm"
   Case "Alpha"
         Page = "\alpha.htm"
   Case "Spatial"
         Page = "\spatial.htm"
   Case "Spatialhier"
         Page = "\SpatialHier.htm"
   Case Else
         Page = "\alpha.htm"
End Select
WebBrowser1.Navigate App.Path & Page
End Sub
```

Appendix C

Listing # 3 – Subroutine #5, #6 Browser Form Code (Operational)

```
Private Sub WebBrowser1_DocumentComplete(ByVal pDisp As Object, URL As Variant)
'*************************************************************
'Event: User navigation to new document page
'Response:
'   Log time taken on last page
'   Prepare for next log entry by recording the page navigated to
'     as the last page visited
'*************************************************************
LogTime
Let LastPage = URL
End Sub
Private Sub LogTime()
'*************************************************************
'Event: Called directly by coded routines
'Response:
'   Get the time spent on the last page and record it in the log
'   Reset the start time for the next page
'*************************************************************
Let EndTime = Timer
Let PageTime = EndTime - StartTime
Write #2, SubjectInfo.Name, LastPage, PageTime
Let StartTime = Timer
End Sub
```

Appendix D

Listing # 3 – Subroutine #3 Browser Form Code (Operational)

```
Private Sub Found_Click()
'********************************************************************
'Event: Found answer button is clicked
'Response:
'  If attempting to select answer from the navigation aid
'      Indicate this action is in error and continue current question
'  Otherwise
'      Log time taken on last page
'      Set up for user to submit the answer
'********************************************************************
If (("\" & WebBrowser1.LocationName) = Page) Then
   MsgBox ("Answers cannot be found on this page.")
Else
  LogTime
  Found.Visible = False
  Submit.Visible = True
End If
End Sub
```

Appendix E

Listing # 3 – Subroutine #4 Browser Form Code (Operational)

```
Private Sub Submit_Click()
'**********************************************************************
'Event: Submit answer button is clicked
'Response:
'  Empty the clipboard
'  Copy the user highlighted information from the document
'       to the Answer variable by having the  WebBrowser's document object
'       use it's ExecCommand to perform a standard command, in this case copy
'  If there is nothing in Answer then
'     Indicate to the user there has been a mistake and exit
'  Otherwise
'     Write the information to the log files
'     If that was the last question then
'        Log the time for this last event
'        Move to title recall form
'     Otherwise
'        Change to the next question
'        Set up to allow user to select found answer
'        Reset new page start time
'        Navigate to navigation aid
'**********************************************************************
Clipboard.SetText ("")
WebBrowser1.Document.ExecCommand ("Copy")
Let Answer = Clipboard.GetText
Let Done = "Done"
If (Answer = "") Then
   MsgBox ("Please highlight the answer before clicking submit")
Else
   Write #3, SubjectInfo.Name, Question, Answer
   Write #2, Done
...Question = Question + 1
...If (Question = 11) Then
          LogTime
          Me.Hide
          TitleRecall.Show
  Else
          Change_Question (Question)
          Submit.Visible = False
          Found.Visible = True
          WebBrowser1.Navigate App.Path & Page
  End If
End If
StartTime = Timer
End Sub
```